arXiv:1401.1872v1 [cs.DB] 9 Jan 2014

# Skew in Parallel Query Processing

Paul Beame, Paraschos Koutris and Dan Suciu
University of Washington, Seattle, WA
{beame,pkoutris,suciu}@cs.washington.edu

January 10, 2014

**Abstract**

We study the problem of computing a conjunctive query $q$ in parallel, using $p$ of servers, on a large database. We consider algorithms with one round of communication, and study the complexity of the communication. We are especially interested in the case where the data is skewed, which is a major challenge for scalable parallel query processing. We establish a tight connection between the *fractional edge packings* of the query and the *amount of communication*, in two cases. First, in the case when the only statistics on the database are the cardinalities of the input relations, and the data is skew-free, we provide matching upper and lower bounds (up to a poly log $p$ factor) expressed in terms of fractional edge packings of the query $q$. Second, in the case when the relations are skewed and the heavy hitters and their frequencies are known, we provide upper and lower bounds (up to a poly log $p$ factor) expressed in terms of packings of *residual queries* obtained by specializing the query to a heavy hitter. All our lower bounds are expressed in the strongest form, as number of bits needed to be communicated between processors with unlimited computational power. Our results generalizes some prior results on uniform databases (where each relation is a matching) [4], and other lower bounds for the MapReduce model [1].

## 1 Introduction

While in traditional query processing the main complexity is dominated by the disk access time, in modern massively distributed systems the dominant cost is that of the communication. A data analyst will use a cluster with sufficiently many servers to ensure that the entire data fits in main memory. Unlike MapReduce [5] which stores data on disk between the Map and the Reduce phase for recovery purposes, newer systems like Spark [16] and its SQL-extension Shark [13] perform the entire computation in main memory, and use replay to recover. In these systems the new complexity parameter is the communication cost, both the amount of data sent and the number of rounds.

A key requirement in such systems is that the data be uniformly partitioned on all servers, and this requirement is challenging to enforce when the input data is *skewed*. A value in the database is skewed, and is called a *heavy hitter*, when its frequency is much higher than some predefined threshold. Since data reshuffling is typically done using hash-partitioning, all records containing a heavy hitter will be sent to the same server, causing it to be overloaded. Skew for parallel joins

has been studied intensively since the early days of parallel databases, see [12]. The standard join algorithm that handles skew is the *skew join* [10] and consists of first detecting the heavy hitters (e.g. using sampling), then treating them differently from the others values, e.g. by partitioning tuples with heavy hitters on the other attributes; a detailed description is in [14]. None of these algorithms has been proven to be optimal in any formal sense, and, in fact, there are no lower bounds for the communication required to compute a join in the presence of skew.

Complex queries often involve multiple joins, and the traditional approach is to compute one join at a time leading to a number of communication rounds at least as large as the depth of the query plan. It is possible, however, to compute a multiway join in a single communication round, using a technique that can be traced back to Ganguli, Silberschatz, and Tsur [7, Sec.7], and was described by Afrati and Ullman [2] in the context of MapReduce algorithms. We will refer to this technique as the HYPERCUBE *algorithm*, following [4]. The $p$ servers are organized into a hypercube with $k$ dimensions, where $k$ is the number of variables in the query. During a single reshuffling step, every tuple is sent to all servers in a certain subcube of the hypercube (with as many dimensions as the number of query variables missing from the tuple). One challenge in this approach is to determine the size of the hypercube in each of the $k$ dimensions. In [2] this is treated as a non-linear optimization problem and solved using Lagrange multipliers. In [4] it is shown that, in the case where all relations have the same cardinality, the optimal dimensions are expressed in terms of the optimal fractional vertex cover of the query. All hypercube-based techniques described in [2, 4], and elsewhere (e.g. in [11] for computing triangles) assume that the data has no skew. The behavior of the algorithm on skewed data has not been studied before and no techniques for addressing skew have been proposed.

**Our Contribution.** In this paper we study the problem of computing a full conjunctive query (multi-way join query), assuming that the input data may have arbitrary skew. We are given $p$ servers that have to compute a query on a databases with $m$ tuples; we assume $m \gg p$. We prove matching upper and lower bounds (up to poly-log $p$ factor) for the amount of communication needed to compute the query in one communication round. We assume the following statistics on the input database to be known: the cardinality of each input relation, the identity of the heavy hitters, and their (approximate) frequency in the data. We note that this is a reasonable assumption in today's distributed query engines. In our settings there are at most $O(p)$ heavy hitters because we choose a threshold for the frequency of heavy hitters that is $\geq m/p$, and therefore the number of heavy hitters is tiny compared to the size of the database. We assume that at the beginning of the computation all servers know the identity of all heavy hitters, and the (approximate) frequency of each heavy hitter. Given these statistics, we describe an explicit formula and prove that it is both an upper (up to a poly-log $p$ factor) and a lower bound for the amount of communication needed to compute the query on the class of databases satisfying those statistics. Our results are significant extensions of our previous results [4] which hold in the absence of skew.

Grohe and Marx [8] and Atserias et al. [3] give upper bounds on the query size in terms of a fractional *edge cover*; this is also a lower bound on the running time of any sequential algorithm that computes the query. Recently, Ngo at al. [9] described a sequential algorithm that matches that bound. Thus, the sequential complexity of a query is captured by the edge cover; our results show that the communication complexity for parallel evaluation is captured by the *edge packing*.

**Overview of the results.** Our analysis of skew starts with an analysis of skew-free databases,

but with unequal cardinalities. Consider a simple cartesian product, $q(x,y) = S_1(x), S_2(y)$, of two relations with cardinalities $m_1, m_2$. Assume $1/p \leq m_1/m_2 \leq p$[1]. Let $p_1 = \sqrt{m_1 p/m_2}$, $p_2 = \sqrt{m_2 p/m_1}$ and assume they are integer values. Organize the $p$ servers into a $p_1 \times p_2$ rectangle, and assign to each server two coordinates $(i,j) \in [p_1] \times [p_2]$. During the shuffle phase the algorithm uses two hash functions $h_1, h_2$ and sends every tuple $S_1(x)$ to all servers with coordinates $(h_1(x), j)$, $j \in [p_2]$ (thus, every server receives $m_1/p_1 = \sqrt{m_1 m_2/p}$ tuples from $S_1$), and sends every tuple $S_2(y)$ to all servers with coordinates $(i, h_2(y))$, $i \in [p_1]$. The load per server is $L = 2\sqrt{m_1 m_2/p}$, and it is not hard to see that this is optimal[2]. This observation generalises to any $u$-way cartesian product: the minimum load per server needed to compute $S_1 \times \ldots \times S_u$ is $\Omega((m_1 m_2 \cdots m_u/p)^{1/u})$

Consider now some arbitrary full conjunctive query $q$ over relations $S_1, \ldots, S_\ell$, and assume that the cardinality of $S_j$ is $m_j$. Choose some subset $S_{j_1}, S_{j_2}, \ldots, S_{j_u}$; the subset is called an *edge packing*, or an *edge matching*, if no two relations share a common variable. Any one-round algorithm that computes the query correctly must also compute the cartesian product of the relations $S_{j_1}, S_{j_2}, \ldots$ Indeed, since no two relations share variables, any tuple in their cartesian product could potentially be part of the query answer; without knowing the content of the other relations, the input servers that store (fragments of) $S_{j_1}, S_{j_2}, \ldots$ must ensure that any combination reaches some output server. Therefore, the load per server of any one-round algorithm is at least $\Omega((m_{j_1} \cdots m_{j_u}/p)^{1/u})$. Thus, every edge packing gives a lower bound for computing $q$. For example the load per server needed to compute $q(x,y,z,w) = S_1(x,y), S_2(y,z), S_3(z,w)$ is at least $L \geq \sqrt{m_1 m_3/p}$, because of the packing $\{S_1, S_3\}$; it must also be $L \geq m_2/p$, because of the packing $\{S_2\}$. We prove in this paper that this property extends to any *fractional edge packing*. Denote $M_j$ the number of bits needed to represent the relation $S_j$. We show:

**Theorem 1.1.** *Let $\mathbf{u} = (u_1, \ldots, u_\ell)$ be any fractional edge packing for the query $q$, and $u = \sum_j u_j$. Denote $K(\mathbf{u}, \mathbf{M}) = \prod_j M_j^{u_j}$ and $L(\mathbf{u}, \mathbf{M}, p) = (K(\mathbf{u}, \mathbf{M})/p)^{1/u}$. If an algorithm computes $q$ in one step, then at least one server has a load $\Omega(L(\mathbf{u}, \mathbf{M}, p))$. Conversely, let $L_{\texttt{lower}} = \max_{\mathbf{u}} L(\mathbf{u}, \mathbf{M}, p)$ be the maximum over all fractional edge packings. Then, there exists a randomized algorithm for $q$ ( the* HYPERCUBE *algorithm, HC) whose maximum load per server is $O(L_{\texttt{lower}} \ln^k p)$ with high probability on all database without skew.*

In the case when all relations have the same size $M$, then the lower bound is $L(\mathbf{u}, \mathbf{M}, p) = M/p^{1/u}$, whose maximum value is $M/p^{1/\tau^*}$ where $\tau^*$ is the value of the maximal fractional edge packing, and is equal to the fractional vertex covering number for $q$; thus, we recover our prior result in [4], which was stated for the special case when all relations are *matchings*.

Theorem 1.1 completes the analysis of the HC algorithm on skew-free databases with arbitrary cardinalities. In addition, we prove a rather surprising result: the HC algorithm is resilient to skew, in the sense that, even on skewed databases, it can still offer a non-trivial upper bound for the maximum load per server: namely $L = O(m/p^{1/k})$, where $m$ is the largest cardinality, and $k$ the total number of variables in the query. For example, using HC one may compute the join of two relations and guarantee a load of $O(m/p^{1/3})$, *even without any knowledge about skew or*

[1] if $m_1 < m_2/p$ then we can broadcast $S_1$ to all servers and compute the query with a load increase of at most $m_2/p$ per server thus at most double that of any algorithm, because $m_2/p$ is the load required to store $S_2$.

[2] Let $a_i, b_i$ be the number of $S_1$-tuples and $S_2$-tuples received by server $i \in [p]$. On one hand $\sum_i a_i b_i = \langle \bar{a}, \bar{b} \rangle \geq m_1 m_2$ because the servers must report all $m_1 m_2$ tuples; on the other hand $\langle \bar{a}, \bar{b} \rangle \leq \|\bar{a} + \bar{b}\|_2^2/4 \leq p\|\bar{a} + \bar{b}\|_\infty^2/4 = pL^2/4$.

*heavy hitters*. In contrast, a standard hash-join algorithm may incur a load of $\Omega(m)$ when the join attributes have a single value.

Next, we consider the case when information about heavy hitters is known. In addition to knowing the cardinalities of the input relations, we assume that the identities of the heavy hitters are known, and that the frequency in the data of every heavy hitter is also known. For example, if the relation $S_j$ contains an attribute $x$, then we assume to know the set of heavy hitters $H$, together with the frequencies $m_j(h) = |\sigma_{x=h}(S_j)|$, which, by definition, are $\geq m_j/p$.

For this setting we generalize the results for skew-free by proving a lower bound, and a matching upper bound. Our lower bound is an elegant generalization of that in Theorem 1.1, and is expressed in terms of fractional edge packings of *residual queries*: for each set of variables $\mathbf{x}$, the residual query $q_{\mathbf{x}}$ is obtained from $q$ by simply removing the variables $\mathbf{x}$. Our matching upper bound is based on the idea of running the main query on the subset of the database that consists of light hitters, then handling each heavy hitter separately, by computing a residual query. The algorithm is difficult, because of two challenges. First, one needs to consider sets of attributes of each relation $S_j$ that may be heavy hitters jointly, even if none of them is a heavy hitter by itself. Second, an attribute value may become a heavy hitter in the residual query even though it was light in the main query. Our algorithm addresses these challenges by creating, for each subset of attributes of each relation, $O(\log p)$ bins of heavy hitters, where all heavy hitters in a bin have frequencies that differ by at most a factor of two. (Because of this it suffices for our algorithm to have access only to approximate frequencies of heavy hitters.) By considering separately all combinations of bins, we can run residual queries on databases where the frequencies are guaranteed to be uniform, thus avoiding the difficulties that arise from recursion. Denote $M_j(h)$ the number of bits needed to represent the subset $\sigma_{x=h}(S_j)$ of $S_j$. Our second main result is:

**Theorem 1.2.** *Consider all database instances defined by a set of statistics consisting of the cardinalities of the relations, the set of heavy hitters, and the frequency of each heavy hitter. For a set of variables* $\mathbf{x}$ *and any packing* $\mathbf{u}$ *of the residual query* $q_{\mathbf{x}}$ *that saturates the variables in* $\mathbf{x}$*, let* $L_{\mathbf{x}}(\mathbf{u}, \mathbf{M}, p) = \left(\sum_{\mathbf{h}} K(\mathbf{u}, \mathbf{M}(\mathbf{h}))/p\right)^{1/u}$*: then any deterministic algorithm that computes* $q$ *on these databases must have a load* $\geq L_{\mathbf{x}}(\mathbf{u}, \mathbf{M}, p)$*. Moreover, denoting* $L_{lower} = \max_{\mathbf{x},\mathbf{u}} L_{\mathbf{x}}(\mathbf{u}, \mathbf{M}, p)$*, there exists a randomized algorithm for computing* $q$ *whose load per server is* $O(L_{lower} \log^{O(1)} p)$ *with high probability.*

As a final contribution of our paper, we discuss the connection between our results in the MPC model and the results of [1] on models for computation in MapReduce. We show that our results provide new upper and lower bounds for computing conjunctive queries in [1], and in a stronger computational model.

The paper is organized as follows. We give the background in section 3, then present in section 3 our results for the case when the statistics known about the database are restricted to cardinalities. The case of databases with known heavy hitters is discussed in section 4. We present the connection with [1] in section 5 and finally conclude in section 6. Several proofs are relegated to the appendix.

## 2 Preliminaries

We review here the basic definitions from [4].

## 2.1 Massively Parallel Communication

We define here the MPC model. The computation is performed by $p$ servers, called *workers*, connected by a complete network of private[3] channels. The input data is initially distributed evenly among the $p$ workers. The computation proceeds in rounds, where each round consists of local computation at the workers interleaved with global communication. The servers have unlimited computational power, but may be limited in the amount of bits they receive. In this paper, we discuss query evaluation in this model, and consider a single round of communication. The *load* of a server is the number of bits received by the server during the communication; we write $L$ for the maximum load among all servers.

If the size of the input is $M$ bits, an ideal algorithm would split the data equally among the servers, and so we would like to have $L = O(M/p)$; in this case, the total amount of data communicated is $O(M)$ and thus there is no *replication*. Depending on the query, $L$ is higher than the ideal $M/p$ by some factor called *replication factor*. In [4] we considered the case when the input data is perfectly uniform and showed that the replication factor for any conjunctive query is $O(p^\varepsilon)$, where $0 < \varepsilon \leq 1$ is a constant that depends only on the query. In this work we consider arbitrary input data, and the replication factor will be a more complex formula that depends on the database statistics.

**Randomization.** The MPC model allows randomization during the computation. The random bits are available to all servers, and are computed independently of the input data.

**Random Instances and Yao's Lemma.** Our lower bounds are stated by showing that, if the database instance is chosen at random from some known probability space, then any algorithm with a load less than a certain bound can report only $o(1)$ fraction of the expected number of answers to the query. Using Yao's lemma [15] this implies than for any randomized algorithm there exists an instance on which the algorithm will fail with high probability; we refer to [4] for details.

**Input Servers.** In our upper bounds we assume that the input relations $S_j$ are initially partitioned uniformly on the servers: all our algorithm treat tuples in $S_j$ independently of other tuples. For our lower bounds, we assume a more powerful model, where, at the beginning of the algorithm, each relation $S_j$ is stored on a separate server, called an *input server*, which can examine the entire relation in order to determine what message to send. These assumptions are the same as in [4].

**Database Statistics.** In this paper we assume that all input servers know certain database statistics. *Simple database statistics* consists of the cardinalities $m_j$ of all input relations $S_j$; we discuss this case in section 3. *Complex database statistics* add information about heavy hitters; we discuss these in the rest of the paper. The size of these statistics is $O(1)$ in the first case, and $O(p)$ in the second. Both upper and lower bounds assume that these statistics are available to all input servers.

## 2.2 Conjunctive Queries

We consider computing answers to conjunctive queries over an input database in the MPC model. We fix an input vocabulary $S_1, \ldots, S_\ell$, where each relation $S_j$ has arity $a_j$; let $a = \sum_{j=1}^{\ell} a_j$. The

[3] "Private" means that when server $i$ sends a message to server $j$ no other server sees its content.

input data consists of one relation instance for each symbol. We consider full conjunctive queries without self-joins:

$$q(x_1, \ldots, x_k) = S_1(\bar{x}_1), \ldots, S_\ell(\bar{x}_\ell) \tag{1}$$

The query is *full*, meaning that every variable in the body appears the head (for example $q(x) = S(x, y)$ is not full), and *without self-joins*, meaning that each relation name $S_j$ appears only once (for example $q(x, y, z) = S(x, y), S(y, z)$ has a self-join). The *hypergraph* of a query $q$ is defined by introducing one node for each variable in the body and one hyperedge for each set of variables that occur in a single atom. With some abuse we write $i \in S_j$ to mean that the variable $x_i$ occurs in the the variables $vars(S_j)$ of the atom $S_j$.

**Fractional Edge Packing.** A *fractional edge packing* (also known as a *fractional matching*) of a query $q$ is any feasible solution $\mathbf{u} = (u_1, \ldots, u_\ell)$ of the following linear constraints:

$$\begin{aligned} &\forall i \in [k] : \sum_{j: i \in S_j} u_j \leq 1 \\ &\forall j \in [\ell] : u_j \geq 0 \end{aligned} \tag{2}$$

The edge packing associates a non-negative weight $u_j$ to each atom $S_j$ s.t. for every variable $x_i$, the sum of the weights for the atoms that contain $x_i$ do not exceed 1. If all inequalities are satisfied as equalities by a solution to the LP, we say that the solution is *tight*.

For a simple example, an edge packing of the query $L_3 = S_1(x_1, x_2), S_2(x_2, x_3), S_3(x_3, x_4)$ is any solution to $u_1 \leq 1$, $u_1 + u_2 \leq 1$, $u_2 + u_3 \leq 1$ and $u_3 \leq 1$. In particular, the solution $(1, 0, 1)$ is a tight and feasible edge packing. A *fractional edge cover* is a feasible solution $\mathbf{u} = (u_1, \ldots, u_\ell)$ to the system above where $\leq$ is replaced by $\geq$ in Eq.2. Every tight fractional edge packing is a tight fractional edge cover, and vice versa.

### 2.3 Friedgut's Inequality

Friedgut [6] introduces the following class of inequalities. Each inequality is described by a hypergraph, which in our paper corresponds to a query, so we will describe the inequality using query terminology. Fix a query $q$ as in (1), and let $n > 0$. For every atom $S_j(\bar{x}_j)$ of arity $a_j$, we introduce a set of $n^{a_j}$ variables $w_j(\mathbf{a}_j) \geq 0$, where $\mathbf{a}_j \in [n]^{a_j}$. If $\mathbf{a} \in [n]^k$, we denote by $\mathbf{a}_j$ the vector of size $a_j$ that results from projecting on the variables of the relation $S_j$. Let $\mathbf{u} = (u_1, \ldots, u_\ell)$ be a fractional *edge cover* for $q$. Then:

$$\sum_{\mathbf{a} \in [n]^k} \prod_{j=1}^{\ell} w_j(\mathbf{a}_j) \leq \prod_{j=1}^{\ell} \left( \sum_{\mathbf{a}_j \in [n]^{a_j}} w_j(\mathbf{a}_j)^{1/u_j} \right)^{u_j} \tag{3}$$

We illustrate Friedgut's inequality on $C_3$:

$$C_3(x, y, z) = S_1(x, y), S_2(y, z), S_3(z, x) \tag{4}$$

$C_3$ has cover $(1/2, 1/2, 1/2)$. Thus, we obtain the following, where $a, b, c$ stand for $w_1, w_2, w_3$ respectively:

$$\sum_{x,y,z \in [n]} a_{xy} \cdot b_{yz} \cdot c_{zx} \leq \sqrt{\sum_{x,y \in [n]} a_{xy}^2 \sum_{y,z \in [n]} b_{yz}^2 \sum_{z,x \in [n]} c_{zx}^2}$$

Friedgut's inequalities immediately imply a well known result developed in a series of papers [8, 3, 9] that gives an upper bound on the size of a query answer as a function on the cardinality of the relations. For example in the case of $C_3$, consider an instance $S_1, S_2, S_3$, and set $a_{xy} = 1$ if $(x, y) \in S_1$, otherwise $a_{xy} = 0$ (and similarly for $b_{yz}, c_{zx}$). We obtain then $|C_3| \leq \sqrt{|S_1| \cdot |S_2| \cdot |S_3|}$.

# 3 Simple Database Statistics

In this section we consider the case when the statistics on database consists of the cardinalities $m_1, \ldots, m_\ell$ of the relations $S_1, \ldots, S_\ell$. All input servers know these statistics. We denote $\mathbf{m} = (m_1, \ldots, m_\ell)$ the vector of cardinalities, and $\mathbf{M} = (M_1, \ldots, M_\ell)$ the vector of the sizes expressed in bits, where $M_j = a_j m_j \log n$, and $n$ is the size of the domain of each attribute.

## 3.1 The HyperCube Algorithm

We present here the HYPERCUBE (HC) algorithm and its analysis.

The HC algorithm, first described in [2], expresses the number of servers $p$ as $p = p_1 \cdot p_2 \cdots p_k$, where each $p_i$ is called the *share* for the variable $x_i$. The algorithm uses $k$ independently chosen random hash functions $h_i : [n] \to [p_i]$, one for each variable $x_i$. During the communication step, the algorithm sends every tuple $S_j(\mathbf{a}_j) = S_j(a_{i_1}, \ldots, a_{i_{r_j}})$ to all servers $\mathbf{y} \in [p_1] \times \cdots \times [p_k]$ such that $h_{i_m}(a_{i_m}) = \mathbf{y}_{i_m}$ for any $1 \leq m \leq r_j$. In other words, for every tuple in $S_j$, after applying the hash functions the algorithm knows the coordinates for the dimensions $i_1, \ldots, i_{r_j}$ in the hypercube, but does not know the other coordinates, and it simply replicates the tuple along those other dimensions. The algorithm finds all answers, because each potential output tuple $(a_1, \ldots, a_k)$ is known by the server $\mathbf{y} = (h_1(a_1), \ldots, h_k(a_k))$.

Since the HC algorithm is parametrized by the choice of shares, we next address two issues. First, we choose the shares $p_i$ so as to minimize the expected load per server. Second, we prove that, with high probability on the choices of the random hash functions, the expected load is not exceeded by more than a factor for any server. We start with the latter, which was not addressed in [2], and was addressed only in a limited setting in [4]: our analysis reveals a previously unknown property of the HC algorithm.

**Analysis of the Load Per Server.** Our analysis is based on the following lemma about hashing.

**Lemma 3.1.** *Let $R(A_1, \ldots, A_r)$ be a relation of arity $r$ with at most $m$ tuples. Let $p_1, \ldots, p_r$ be integers and denote $p = \prod_i p_i$ where $m \geq p$. Suppose that we hash each tuple $(a_1, \ldots, a_r)$ to the bucket $(h_1(a_1), \ldots, h_r(a_r))$, where $h_1, \ldots, h_r$ are independent and perfectly random hash functions. Then:*

1. *The expected load in every bucket is $m/p$.*
2. *If for every $i \in [r]$, every value of the attribute $A_i$ occurs at most once, then the maximum load per bucket is $O(m/p)$ with high probability.*
3. *If for every $S \subseteq [r]$, every tuple of values of attributes $(A_i)_{i \in S}$ occurs at most $am/\prod_{i \in S} p_i$ times, then the maximum load per bucket is $O(a^r \ln^r(p) m/p)$ with high probability.*
4. *The maximum load per bucket is $O(m/\min_i(p_i))$, independent of the instance.*

We prove this lemma in Appendix B. For the case where the arity of the relation is $r = 1$, the above lemma is a straightforward application of Chernoff bounds. The case where $r \geq 2$ requires a more sophisticated argument and is novel, to the best of our knowledge.

We can apply the lemma to analyze the behavior of the HC algorithm under two conditions: over skew-free databases, and over arbitrary databases. Given a vector of shares $(p_1, \ldots, p_k)$, we say that a relation $S_j$ is *skew-free* w.r.t. the shares if for every subset of variables $\mathbf{x} \subseteq vars(S_j)$, every value has frequency at most $m_j / \prod_{x_i \in \mathbf{x}} p_i$. Our prior analysis in [4] was only for the special case when the frequency of each value at each attribute is at most 1.

**Corollary 3.2.** *Let $\mathbf{p} = (p_1, \ldots, p_k)$ be the shares of the HC algorithm.*

*(i) If $S_j$ is skew-free w.r.t. $\mathbf{p}$, then with high probability the maximum load per server is*

$$O\left(\max_j \frac{M_j}{\prod_{i:i \in S_j} p_i} \ln^k(p)\right)$$

*(ii) For any given database, with high probability the maximum load per server is*

$$O\left(\max_j \frac{M_j}{\min_{i:i \in S_j}(p_i)}\right)$$

In [2], it is assumed that the database is skew-free and that the load per server is the expected load; item (i) of our result confirms that the load does not exceed the expected load by more than a poly-log factor with high probability (our proof uses Chernoff bounds). Item (ii) is novel, because it describes how the HC algorithm behaves on skewed data: it shows that the algorithm is resilient to skew, and gives an upper bound even on skewed databases. We illustrate with an example.

**Example 3.3.** *Let $q(x, y, z) = S_1(x, z), S_2(y, z)$ be a simple join, where both relations have cardinality $m$. We show two instances of the HC algorithm, the first optimized for skewed databases, and the second optimized for skew-free databases. The first share allocation is $p_1 = p_2 = p_3 = p^{1/3}$, thus every processor is identified by $(w_1, w_2, w_3) \in [p_1] \times [p_2] \times [p_3]$. The algorithm sends every tuple $S_1(a, c)$ to all processors $(h_1(a), w_2, h_3(c))$ for $w_2 \in [p_3]$ and every tuple $S_2(b, c)$ to all processors $(w_1, h_2(b), h_3(c))$ for $w_1 \in [p_1]$. By Corollary 3.2, on skew-free databases the load per server is $O(m/p^{2/3})$ (plus polylog factor). But even on skewed database the load per server is $O(m/p^{1/3})$. The second algorithm allocates shares $p_1 = p_2 = 1, p_3 = p$. This corresponds to a standard hash-join on the variable $z$. On a skew-free database (equivalently, when every value of $z$ has frequency at most $m/p$ in both relations) the load per server is $O(m/p)$ with high probability. However, if it is skewed, then the load can be as bad as $O(m)$: this occurs when all tuples have the same value $z$.*

Generalizing the example, for every conjunctive query with $k$ variables, we can execute the HC algorithm with equal shares $p_1 = \ldots = p_k = p^{1/k}$. Then, the algorithm achieves a maximum load per server of at most $O(\max_j M_j / p^{1/k})$.

However, in practice, in applications where skew is expected, it is better to design specialized algorithms, as we further discuss in section 4. Therefore, we focus our analysis on skew-free databases, and optimize the expected load.

**Choosing the Shares.** Here we discuss how to compute the shares $p_i$ to optimize the expected load per server. Afrati and Ullman compute the shares by optimizing the total load $\sum_j m_j / \prod_{i:i \in S_j} p_i$

subject to the constraint $\prod_i p_i = 1$, which is a non-linear system that can be solved using Lagrange multipliers. Here we take a different approach. First, we write the shares as $p_i = p^{e_i}$ where $e_i \in [0,1]$ is called the *share exponent* for $x_i$, and denote $L$ the maximum load per server, thus $M_j / \prod_{i:i\in S_j} p_i \le L$ for every $j$. Denote $\lambda = \log_p L$ and $\mu_j = \log_p M_j$ (we will assume w.l.o.g. that $m_j \ge p$, hence $\mu_j \ge 1$ for all $j$). Then, we optimize the LP:

$$\begin{aligned} \text{minimize} \quad & \lambda \\ \text{subject to} \quad & \sum_{i\in[k]} -e_i \ge -1 \\ & \forall j \in [\ell] : \sum_{i\in S_j} e_i + \lambda \ge \mu_j \\ & \forall i \in [k] : e_i \ge 0, \quad \lambda \ge 0 \end{aligned} \tag{5}$$

Denote $L_{\texttt{upper}} = p^{e^*}$ where $e^*$ is the objective value of the optimal solution to the above LP. We have:

**Theorem 3.4.** *For a query $q$ and $p$ servers, with statistics $\mathbf{M}$, let $\mathbf{e} = e_1, \ldots, e_k$ be share exponents that are optimal for the above LP. Then, the expected load per server is $L_{\texttt{upper}}$. Moreover, if every $S_j$ is skew-free w.r.t. to $\mathbf{e}$, then the maximum load per server is $O(L_{\texttt{upper}} \cdot \ln^k(p))$ with high probability.*

In subsection 3.3 we will give a closed form expression for $L_{\texttt{upper}}$ and also provide an example. But first, we prove a matching lower bound.

### 3.2 The Lower Bound

We next prove a lower bound for the maximum load per server over databases with statistics $\mathbf{M}$. Fix some constant $0 < \delta < \min_j\{a_j\}$, and assume that for every relation $S_j$, its cardinality satisfies $m_j \le n^\delta$, where $n$ is the domain size of each attribute.

Consider the probability space where each relation $S_j$ is chosen independently and uniformly at random from all subsets of $[n]^{a_j}$ with exactly $m_j$ tuples. Denote $\mathbf{E}[|q(I)|]$ the expected number of answers to $q$. (We show in the appendix, Lemma A.1, that $\mathbf{E}[|q(I)|] = n^{k-a}\prod_{j=1}^{\ell} m_j$.)

Fix a query $q$ and a fractional edge packing $\mathbf{u}$ of $q$. Denote $u = \sum_{j=1}^{\ell} u_j$ the value of the packing, and:

$$K(\mathbf{u}, \mathbf{M}) = \prod_{j=1}^{\ell} M_j^{u_j} \tag{6}$$

$$L(\mathbf{u}, \mathbf{M}, p) = \left(\frac{K(\mathbf{u}, \mathbf{M})}{p}\right)^{1/u} \tag{7}$$

Further denote $L_{\texttt{lower}} = \max_{\mathbf{u}} L(\mathbf{u}, \mathbf{M}, p)$, where $\mathbf{u}$ ranges over all edge packings for $q$. Let $c$ be a constant, $c = \frac{a_j - \delta}{3a_j}$, where $a_j$ is the maximum arity of all relations. We prove in Appendix A:

**Theorem 3.5.** *Fix statistics $\mathbf{M}$, and consider any deterministic MPC algorithm that runs in one communication round on $p$ servers. (1) Let $\mathbf{u}$ be any edge packing of $q$. If $i$ is any server and $L_i$ is its load, then server $i$ reports at most*

$$\frac{L_i^u}{c^u K(\mathbf{u}, \mathbf{M})} \mathbf{E}[|q(I)|]$$

*answers in expectation, where $I$ is a randomly chosen database with statistics* $\mathbf{M}$*. Therefore, the $p$ servers of the algorithm report at most*

$$\left(\frac{L}{c \cdot L(\mathbf{u}, \mathbf{M}, p)}\right)^u \cdot \mathbf{E}[|q(I)|]$$

*answers in expectation, where $L$ is the maximum load of all servers.*

*As a consequence, any algorithm that computes $q$ correctly over databases with statistics $M$ must have load $L \geq cL_{lower}$ bits*[4]*.*

In our previous work [4], we presented a matching lower/upper bound for computing $q$ on some restricted database instances, where the relations $S_j$ are *matchings* and have the same cardinalities; the proof of Theorem 3.5 is an extension of the lower bound proof in [4]. We explain here the relationship. When all cardinalities are equal, $M_1 = \ldots = M_\ell = M$, then $L_{\texttt{lower}} = M/p^{1/u}$, and this quantity is maximized when $\mathbf{u}$ is a maximum fractional edge packing, whose value is denoted $\tau^*$: this is equal to the fractional vertex covering number for $q$. The bound in [4] is $M/p^{1/\tau^*}$. (The constant $c$ in [4] is tighter.) Theorem 3.5 generalizes the lower bound to arbitrary cardinalities and, in that case, $L(\mathbf{u}, \mathbf{M}, p)$ is not necessarily maximized at $\tau^*$.

In the rest of this section we prove that $L_{\texttt{lower}} = L_{\texttt{upper}}$ and also give a closed form expression for both.

### 3.3 Proof of Equivalence

If $\mathbf{u}, \mathbf{u}'$ are two fractional edge packings for a conjunctive query $q$, then we write $\mathbf{u} \leq \mathbf{u}'$ if $u_j \leq u'_j$ for all $j$ and say that $\mathbf{u}'$ *dominates* $\mathbf{u}$. Let $pk(q)$ be the non-dominated vertices of the polytope defined by the fractional edge packing constraints in (2). Recall that the vertices of the polytope are feasible solutions $\mathbf{u}_1, \mathbf{u}_2, \ldots$, with the property that every other feasible solution $\mathbf{u}$ to the LP is a convex combination of these vertices. Each vertex can be obtained by choosing $m$ out of the $k + \ell$ inequalities in (2), transforming them into equalities, then solving for $\mathbf{u}$: $pk(q)$ contains the subset of vertices that are not dominated by others, and $|pk(q)| \leq \binom{k+\ell}{m}$. We prove here:

**Theorem 3.6.** *For any vector of statistics* $\mathbf{M}$ *and number of processors $p$ , we have:*

$$L_{lower} = L_{upper} = \max_{\mathbf{u} \in pk(q)} L(\mathbf{u}, \mathbf{M}, p)$$

Before we prove the theorem we discuss its implications. We start with an example.

**Example 3.7.** *Consider the triangle query*

$$C_3 = S_1(x_1, x_2), S_2(x_2, x_3), S_3(x_3, x_1)$$

*and assume the three cardinalities are $m_1, m_2, m_3$. Then, $pk(C_3)$ has four vertices, and each gives a different value for $L(\mathbf{u}, \mathbf{M}, p)$:*

[4]This follows by observing that, when $L(\mathbf{u}, \mathbf{M}, p)$ is maximized, then $u = \sum_j u_j \geq 1$.

| $\mathbf{u}$ | $L(\mathbf{u}, \mathbf{M}, p)$ |
|---|---|
| $(1/2, 1/2, 1/2)$ | $(M_1 M_2 M_3)^{1/3}/p^{2/3}$ |
| $(1, 0, 0)$ | $M_1/p$ |
| $(0, 1, 0)$ | $M_2/p$ |
| $(0, 0, 1)$ | $M_3/p$ |

*The first vertex is the solution to $u_1 + u_2 = u_1 + u_3 = u_2 + u_3 = 1$; the second the solution to $u_1 + u_2 = 1, u_2 = u_3 = 0$, etc. Thus, the load of the algorithm is the largest of these four quantities, and this is also the lower bound of any algorithm. In other words, the optimal solution to the LP (5) can be given in closed form, as the maximum over four expressions.*

Next we use the theorem to compute the *space exponent*. In [4] we showed that, for every query $q$, the optimal load over databases restricted to matchings is $M/p^{1-\varepsilon}$, where $0 \leq \varepsilon < 1$ is called the space exponent for $q$. Consider now a database with arbitrary statistics, and denote $M = \max_j M_j$. If for some $j$, $M_j \leq M/p$, then an optimal algorithm can simply broadcast $M_j$ to all $p$ servers, and remove $S_j$ from the query: this will increase the load by at most a factor of 2. Thus, we may assume w.l.o.g. that for every $j$, $M_j = M/p^{\nu_j}$ for some $\nu_j \in [0, 1)$. Then $L(\mathbf{u}, \mathbf{M}, p) = M/p^{\sum_j \nu_j u_j + 1/u}$. To obtain the optimal load, one needs to find $\mathbf{u} \in pk(q)$ that minimizes $v = \sum_j \nu_j u_j + 1/(\sum_j u_j)$. Denoting $v^*$ the minimal value, the load is $M/p^{v^*}$. Thus, the space exponent for given statistics is $1 - v^*$.

*Proof.* (of Theorem 3.6) Recall that $L_{\texttt{upper}}$ is $p^{e^*}$, where $e^*$ is the optimal solution to the *primal* LP problem (5). Consider its *dual* LP:

$$\begin{aligned}
\text{maximize} \quad & \sum_{j \in [\ell]} \mu_j f_j - f \\
\text{subject to} \quad & \sum_{j \in [\ell]} f_j \leq 1 \\
\forall i \in [k] : & \sum_{j : i \in S_j} f_j - f \leq 0 \\
\forall j \in [\ell] : & f_j \geq 0, \quad f \geq 0
\end{aligned} \tag{8}$$

By the primal/dual theorem, its optimal solution is also $e^*$. Writing $u_j = f_j/f$ and $u = 1/f$, we transform it into the following non-linear optimization problem:

$$\begin{aligned}
\text{maximize} \quad & \frac{1}{u} \cdot \left( \sum_{j \in [\ell]} \mu_j u_j - 1 \right) \\
\text{subject to} \quad & \sum_{j \in [\ell]} u_j \leq u \\
\forall i \in [k] : & \sum_{j : i \in S_j} u_j \leq 1 \\
\forall j \in [\ell] : & u_j \geq 0
\end{aligned} \tag{9}$$

Consider optimizing the above non-linear problem. Its optimal solution must have $u = \sum_j u_j$, otherwise we simply replace $u$ with $\sum_j u_j$ and obtain a feasible solution with at least as good objective function (indeed, $\mu_j \geq 1$ for any $j$, and hence $\sum_j \mu_j u_j \geq \sum_j u_j \geq 1$, since any optimal $\mathbf{u}$ will have sum at least 1). Therefore, the optimal is given by a fractional edge packing $\mathbf{u}$. Furthermore, for any packing $\mathbf{u}$, the objective function $\sum_j \frac{1}{u} \cdot (\mu_j u_j - 1)$ is $\log_p L(\mathbf{u}, \mathbf{M}, p)$. To prove the theorem, we show that (a) $e^* = u^*$ and (b) the optimum is obtained when $\mathbf{u} \in pk(q)$. This follows from:

**Lemma 3.8.** *Consider the function $F : \mathbf{R}^{k+1} \to \mathbf{R}^{k+1}$: $F(x_0, x_1, \ldots, x_k) = (1/x_0, x_1/x_0, \ldots, x_k/x_0)$. Then:*

- *$F$ is its own inverse, $F = F^{-1}$.*
- *$F$ maps any feasible solution to the system* (8) *to a feasible solution to* (9)*, and conversely.*
- *$F$ maps a convex set to a convex set.*

*Proof.* If $y_0 = 1/x_0$ and $y_j = x_j/x_0$, then obviously $x_0 = 1/y_0$ and $x_j = y_j/y$. The second item can be checked directly. For the third item, it suffices to prove that $F$ maps a convex combination $\lambda\mathbf{x} + \lambda'\mathbf{x}'$ where $\lambda + \lambda' = 1$ into a convex combination $\mu F(\mathbf{x}) + \mu' F(\mathbf{x}')$, where $\mu + \mu' = 1$. Assuming $\mathbf{x} = (x_0, x_1, \ldots, x_k)$ and $\mathbf{x}' = (x'_0, x'_1, \ldots, x'_k)$, this follows by setting $\mu = x_0/(\lambda x_0 + \lambda x'_0)$ and $\mu' = x'_0/(\lambda x_0 + \lambda x'_0)$. □

This completes the proof of Theorem 3.6. □

# 4 Complex Database Statistics

In this section, we discuss algorithms and lower bounds in the case where the input servers are provided by additional information regarding heavy hitters.

## 4.1 A Simple Case: Join

We start with a simple example, the join of two tables, $q(x, y, z) = S_1(x, z), S_2(y, z)$, to illustrate the main algorithmic and proof ideas. The algorithm uses the same principle popular in virtually all parallel join implementations to date: identify the heavy hitters and treat them differently. However, the analysis and optimality proof is new, to the best of our knowledge.

Let $m_1, m_2$ be the cardinalities of $S_1, S_2$. For any value $h \in [n]$ that variable $z$ may assume, let $m_j(h)$ denote the frequency of $h$ in $S_j$, $j = 1, 2$; $h$ is called a *heavy hitter* in $S_j$ if $m_j(h) \geq m_j/p$. Let $H_{12}$ denote the set of all heavy hitters in both $S_1$ and $S_2$, $H_1$ those heavy only in $S_1$, and $H_2$ those heavy only in $S_2$. Note that $|H_{12}|, |H_1|, |H_2| \leq p$. We assume that all heavy hitters and their frequencies are known by the algorithm. A tuple $S_j(a, b)$, for $j = 1, 2$, is called a *light hitter* if the value $b$ is in none of the sets $H_1, H_2, H_{12}$.

For each $h \in H_{12}$, we must compute a cartesian product $q[h/z] = S_1(x, h), S_2(y, h)$: if we allocate $p_h$ servers to this task, then we have seen in Section 1 that the optimal way is to write $p_h = p_1 \cdot p_2$ where $p_1 = \lceil p_h \sqrt{m_1(h)/m_2(h)} \rceil$ and $p_2 = \lceil p_h \sqrt{m_2(h)/m_1(h)} \rceil$, then use the algorithm in Section 1. However, if $h \in H_1$, we cannot apply the same formula because $m_1(h) \gg p_h m_2(h)$

and we may obtain $p_1 \gg p_h$: in that case we allocate all servers to the variable $x$. Define:

$$\forall h \in H_{12}: \quad K_{12}(h) = m_1(h) \cdot m_2(h)$$
$$p_h = \left\lceil p \frac{K_{12}(h)}{\sum_{h' \in H_{12}} K_{12}(h')} \right\rceil, L_{12} = \sqrt{\frac{\sum_{h' \in H_{12}} K_{12}(h')}{p}}$$

and similarly for the other two types $j = 1, 2$:

$$\forall h \in H_j: \quad K_j(h) = m_j(h)$$
$$p_h = \left\lceil p \frac{K_j(h)}{\sum_{h' \in H_j} K_j(h')} \right\rceil, L_j = \sqrt{\frac{\sum_{h' \in H_j} K_j(h')}{p}}$$

The algorithm operates as follows:

1. Compute $S_1(x,z), S_2(y,z)$ using a hash-join on the variable $z$, on all light hitters, using all $p$ servers.
2. For a heavy hitter $h \in H_{12}$: use $p_h$ servers to compute the cartesian product $S_1(x,h), S_2(y,h)$.
3. For a heavy hitter $h \in H_1$: use $p_h$ servers. Hash-partition $S_1(x,h)$ on $x$, and broadcast $S_2(y,h)$ to all $p_h$ servers.
4. For a heavy hitter $h \in H_2$: similar to above.

While our description has four logical steps, the algorithm consists of a single communication step. All input servers know the heavy hitters, and thus can compute all quantities $p_h$. Then, in parallel, every input server holding some fragment of $S_j$ will examine each tuple $S_j(a,b)$ and determine in which case it is handled.

For each of the steps of the algorithm, the number of servers used may exceed $p$ (since we round up to the closest integer), but we can show that it will always be $\Theta(p)$. Additionally, by examining the complexity of each step, we can conclude that, with high probability, the maximum load per server is at most $O(L \ln(p))$,

$$L = \max(m_1/p, m_2/p, L_1, L_2, L_{12}) \tag{10}$$

We end our discussion by showing that the algorithm is almost optimal; we argue next that any algorithm must incur a load per server of at least $L$. It suffices to prove that, for each of the quantities under $\max(\ldots)$ in (10), any algorithm must have a load of at least that quantity. We illustrate here the main idea of the proof for $L_{12}$. Let $[n]$ denote the domain; then the size of the join is $\sum_{h \in [n]} m_1(h) m_2(h)$. Denote $w_{i1}(h)$ the number of tuples of the form $S_1(x,h)$ received by server $i$; $\mathbf{w}_1(h) = (w_{11}(h), \ldots, w_{p1}(h))$ be the $p$-vector representing how the tuples with $z = h$ in $S_1$ are distributed to the $p$ servers; and $\mathbf{w}_1 = \sum_{h \in [n]} \mathbf{w}_1(h)$ the vector showing the total load per server coming from the table $S_1$. We denote similarly $w_{i2}(h)$, $\mathbf{w}_2(h)$, $\mathbf{w}_2$ for $S_2$. Server $i$ can report

at most $w_{i1}(h)w_{i2}(h)$ join tuples with $z = h$. Since the servers must find all answers,

$$\sum_{h\in[n]} m_1(h)m_2(h) \le \sum_{i=1}^{p}\sum_{h\in[n]} w_{i1}(h)w_{i2}(h)$$
$$\le \sum_{i=1}^{p}\sum_{h\in[n]} w_{i1}(h)\cdot \sum_{h\in[n]} w_{i2}(h) = \langle \mathbf{w}_1, \mathbf{w}_2\rangle$$
$$\le \|\mathbf{w}_1\|_2\|\mathbf{w}_2\|_2 \le p\|\mathbf{w}_1\|_\infty\|\mathbf{w}_2\|_\infty \le pL^2$$

In the last line we used the fact that $\|\mathbf{w}\|_2 \le \sqrt{p}\|\mathbf{w}\|_\infty$ and that $L \ge \max(\|\mathbf{w}_1\|_\infty, \|\mathbf{w}_2\|_\infty)$.

### 4.2 An Algorithm for the General Case

We now generalize some of the ideas for the simple join to an arbitrary conjunctive query $q$. Extending the notion for simple joins, for each relation $S_j$ with $|S_j| = m_j$ we say that a partial assignment $\mathbf{h}_j$ to a subset $\mathbf{x}_j \subset vars(S_j)$ is a *heavy hitter* if and only if the number of tuples, $m_j(\mathbf{h}_j)$, from $S_j$ that contain $\mathbf{h}_j$ satisfies $m_j(\mathbf{h}_j) > m_j/p$. As before, there are $O(p)$ such heavy hitters. We will assume that each input server knows the entire set of heavy hitters for all relations.

For simplicity we assume that $p$ is a power of 2. We will not produce quite as smooth a bound as we did for the simple join, but we will show that the bound we produce is within a $\log^{O(1)} p$ factor of the optimal. To do this, for each relation $S_j$ and subset of variables $\mathbf{x}_j$, we define $\log_2 p$ bins for the frequencies, or degrees of each of the heavy hitters. The $b$-th bin, for $b = 1, \ldots, \log_2 p$ will contain all heavy hitters $\mathbf{h}_j$ with $m_j/2^{b-1} \ge m_j(\mathbf{h}_j) > m_j/2^b$. The last bin, a bin of light hitters with $b = \log_2 p + 1$, will contain all assignments $\mathbf{h}_j$ to $\mathbf{x}_j$ that are not heavy hitters. Notice that, when $\mathbf{x}_j = \emptyset$, then the only non-empty bin is the first bin, the only heavy hitter is the empty tuple $\mathbf{h}_j = ()$, and $m_j(\mathbf{h}_j) = m_j$.

For a bin $b$ on $\mathbf{x}_j$ define $\beta_b = \log_p(2^{b-1})$; observe that for each heavy hitter bin, there are at most $2p^{\beta_b}$ heavy hitters in this bin, and for the last bin we have $\beta_b = 1$. Instead of identifying each bin using its index $b$, we identify each bin by $\beta_b$, called its *bin exponent*, along with the index of the relation $S_j$ for which it is defined, and the set $\mathbf{x}_j \subset vars(S_j)$. Note that $0 = \beta_1 < \beta_2 < \cdots < \beta_{\log_2 p+1} = 1$.

**Definition 4.1** (Bin Combination)**.** *Let $\mathbf{x} \subset V = vars(q)$, and define $\mathbf{x}_j = \mathbf{x} \cap vars(S_j)$. A pair $\mathcal{B} = (H, (\beta_j)_j)$ is called a* bin combination *if (1) $\beta_j = 0$ for every $j$ where $\mathbf{x}_j = \emptyset$, and (2) there is some consistent assignment $\mathbf{h}$ to $\mathbf{x}$ such that for each $j$ with $\mathbf{x}_j \ne \emptyset$ the induced assignment $\mathbf{h}_j$ to $\mathbf{x}_j$ has bin exponent $\beta_j$ in relation $S_j$. We write $C(\mathcal{B})$ for the set of all such assignments $\mathbf{h}$.*

Our algorithm allocates $p$ virtual processors to each bin combination and handles associated inputs separately. There are $O(\log p)$ bin choices for each relation and therefore at most $\log^{O(1)} p$ bin combinations in total, so at most $p\log^{O(1)} p$ virtual processors in total. Let $N_{\mathrm{bc}}$ be the number of possible bin combinations. As in the join algorithm (subsection 4.1), within each bin combination we partition the $p$ servers among the heavy hitters, using $p_{\mathbf{h}} = p^{1-\alpha}$ servers for heavy hitter $\mathbf{h}$ (independent of $\mathbf{h}$, since we have ensured complete uniformity within a bin combination). However, we can only process $p^\alpha \le p$ heavy hitters in every bin combination: in general, we may have

$C(\mathcal{B}) > p$, e.g. when $\mathbf{x}$ contains a variable $x_1$ in $S_1$ and a variable $x_2$ in $S_2$, then there may be up to $p \times p$ heavy hitters in this bin combination.

For a bin combination $\mathcal{B}$ and assignment $\mathbf{h} \in C(\mathcal{B})$, write $m_j(\mathbf{h})$ for $m_j(\mathbf{h}_j)$. For each $\mathcal{B}$ we will define below a set $C'(\mathcal{B}) \subseteq C(\mathcal{B})$ with $|C'(\mathcal{B})| \leq p$ and sets $S_j^{(\mathcal{B})} \subseteq S_j$ of tuples for $j \in [\ell]$ that extend $\mathbf{h}_j$ for some $\mathbf{h} \in C'(\mathcal{B})$.

The algorithm for $\mathcal{B}$ will compute all query answers that are joins of $(S_j^{(\mathcal{B})})_j$, by executing the HC algorithm for a particular choice of exponents given. The share exponents for the HC algorithm will be provided by a modification of the vertex covering primal LP (5), which describes an algorithm that suffices for all light hitters. Recall that in this LP, $\mu_j = \log_p M_j$ and $\lambda$ is $\log_p L$ for the load $L$. That LP corresponds to the bin combination $\mathcal{B}_\emptyset$ which has $\mathbf{x} = \emptyset$ and all $\beta_j = 0$. In this case $C'(\mathcal{B}_\emptyset) = C(\mathcal{B}_\emptyset)$ has 1 element, the empty partial assignment. More generally, write $\alpha = \log_p |C'(\mathcal{B})|$, the LP associated with our algorithm for bin combination $\mathcal{B}$ is:

$$\begin{aligned} \text{mininimize} \quad & \lambda \\ \text{subject to} \quad & \\ \forall j \in [\ell]: \quad & \lambda + \sum_{x_i \in vars(S_j) - \mathbf{x}_j} e_i \geq \mu_j - \beta_j \\ & \sum_{i \in V - \mathbf{x}} e_i \leq 1 - \alpha \\ \forall i \in V - \mathbf{x}: \quad & e_i \geq 0, \quad \lambda \geq 0 \end{aligned} \tag{11}$$

Thus far, this is only fully specified for $\mathcal{B} = \mathcal{B}_\emptyset$ since we have not defined $C'(\mathcal{B})$ for other choices of $\mathcal{B}$. We define $C'(\mathcal{B})$ inductively based on optimal solutions $(\lambda^{(\mathcal{B}')}, (e_i^{(\mathcal{B}')})_{i \in V - \mathbf{x}'})$ to the above LP applied to bin combinations $\mathcal{B}'$ with $\mathbf{x}' \subset \mathbf{x}$ as follows. (Such solutions may not be unique but we fix one arbitrarily for each bin combination.)

For $\mathbf{h}' \in C'(\mathcal{B}')$, we say that a heavy hitter $\mathbf{h}_j$ of $S_j$ that is an extension of $\mathbf{h}'_j$ to $\mathbf{x}_j$ is *overweight* for $\mathcal{B}'$ if there are more than $N_{\mathrm{b}c} \cdot m_j / p^{\beta_j + \sum_{i \in \mathbf{x}_j - \mathbf{x}'_j} e_i^{(\mathcal{B}')}}$ elements of $S_j$ consistent with $\mathbf{h}_j$. $C'(\mathcal{B})$ consists of all assignments $\mathbf{h} \in C(\mathcal{B})$ such that there is some $j \in [\ell]$, some bin combination $\mathcal{B}'$ on set $\mathbf{x}' \subset \mathbf{x}$ such that $\mathbf{x} - \mathbf{x}' \subseteq vars(S_j)$, and some $\mathbf{h}' \in C'(\mathcal{B}')$ such that $\mathbf{h}$ is an extension of $\mathbf{h}'$ and $\mathbf{h}_j$ is an overweight heavy hitter of $S_j$ for $\mathcal{B}'$. The following lemma, which is proved in the appendix, shows that $\alpha \leq 1$ and thus that the LP for $\mathcal{B}$ is feasible.

**Lemma 4.2.** *For all bin combinations $\mathcal{B}$, $|C'(\mathcal{B})| \leq p$.*

Let $A_\mathcal{B} \subseteq [\ell]$ be the set of all $j$ such that $\mathbf{x}_j \neq \emptyset$. For each $j \in [\ell] - A_\mathcal{B}$, let $S_j^{(\mathcal{B})}$ consist of all tuples in $S_j$ that do not contain any heavy hitter $\mathbf{h}''_j$ of $S_j$ that is overweight for $\mathcal{B}$. For each $j \in A_\mathcal{B}$, and $\mathbf{h} \in C'(\mathcal{B})$ let $S_j^{(\mathcal{B})}(\mathbf{h})$ consist of all tuples in $S_j$ that contain $\mathbf{h}_j$ on $\mathbf{x}_j$ (with bin exponent $\beta_j$) but do not contain any heavy hitter $\mathbf{h}''_j$ of $S_j$ that is overweight for $\mathcal{B}$ and a proper extension of $\mathbf{h}_j$. $S_j^{(\mathcal{B})}$ will be the union of all $S_j^{(\mathcal{B})}(\mathbf{h})$ for all $\mathbf{h} \in C'(\mathcal{B})$.

For each of the $p^\alpha$ heavy hitter assignments $\mathbf{h} \in C'(\mathcal{B})$ the algorithm uses $p^{1-\alpha}$ virtual processors to compute the join of the subinstances $S_j^{(\mathcal{B})}(\mathbf{h})$ for $j \in [\ell]$. Those processors will be allocated

using the HC algorithm that assigns $p^{e_i^{(\mathcal{B})}}$ shares to each variable $x_i \in V - \mathbf{x}$. It remains to show that the load of the algorithm for $\mathcal{B}$ is within a $\log^{O(1)} p$ factor of $p^{\lambda^{(\mathcal{B})}}$, where $\lambda^{(\mathcal{B})}$ is given by the LP for $\mathcal{B}$.

**Lemma 4.3.** *Let* $\mathbf{h}$ *be an assignment to* $\mathbf{x}$ *that is consistent with bin combination* $\mathcal{B}$*. If we hash each residual relation* $S_j^{(\mathcal{B})}(\mathbf{h})$ *on* $vars(S_j) - \mathbf{x}_j$ *using* $p^{e_i^{(\mathcal{B})}}$ *values for each* $x_i \in vars(S_j) - \mathbf{x}_j$*, each processor receives*

$$O\left((N_{\text{bc}} \cdot \ln p)^{r'} \cdot m_j / p^{\min(\beta_j + \sum_{i \in vars(S_j) - \mathbf{x}_j} e_i^{(\mathcal{B})}, 1)}\right)$$

*values with high probability, where* $r' = \max_j(r_j - |\mathbf{x}_j|)$*.*

*Proof.* For $j \in [\ell] - A_{\mathcal{B}}$, $S_j^{(\mathcal{B})}$ only contains tuples of $S_j$ that are not overweight for $\mathcal{B}$, which means that for every $\mathbf{x}_j'' \subseteq vars(S_j)$ and every heavy hitter assignment $\mathbf{h}''$ to the variables of $\mathbf{x}_j''$, there are at most

$$N_{\text{bc}} \cdot m_j / p^{\beta_j + \sum_{i \in \mathbf{x}_j''} e_i^{(\mathcal{B})}} = N_{\text{bc}} \cdot m_j / p^{\beta_j + \sum_{i \in \mathbf{x}_j'' - \mathbf{x}_j} e_i^{(\mathcal{B})}}$$

elements of $S_j$ consistent with $\mathbf{h}''$. Every other assignment $\mathbf{h}''$ to the variables of $\mathbf{x}_j''$ is a light hitter and therefore is contained in at most $m_j/p$ consistent tuples of $S_j$. For $j \in A_{\mathcal{B}}$, we obtain the same bound, where the only difference is that we need to restrict things to extensions of $\mathbf{h}_j$. This bound gives the smoothness condition on $S_j^{(\mathcal{B})}(h)$ necessary to apply Lemma 3.1 to each relation $S_j^{(\mathcal{B})}(\mathbf{h})$ and yields the claimed result. □

As a corollary, we obtain:

**Corollary 4.4.** *Let* $L_{min} = \max_j(m_j/p)$*. The maximum load of our algorithm using* $p$ *(virtual) processors for* $\mathcal{B}$ *is* $O((N_{\text{bc}} \cdot \ln p)^{r_{\max}} \cdot \max(L_{min}, p^{\lambda}))$ *with high probability, where* $\lambda = \lambda^{(\mathcal{B})}$ *is the optimum of the LP for* $\mathcal{B}$ *and* $r_{\max}$ *is the maximum arity of any* $S_j$*.*

*Proof.* There are $p^{1-\alpha}$ processors allocated to each $\mathbf{h}$ and $p^{\alpha}$ such assignments $\mathbf{h}$ so that the maximum load per $\mathbf{h}$ is also the maximum overall load. Given the presence of the $L_{min}$ term, it suffices to show that the maximum load per $\mathbf{h}$ due to relation $S_j^{(\mathcal{B})}(\mathbf{h})$ is at most $(\ln p)^{k-|\mathbf{x}|} \cdot \max(m_j/p, p^{\lambda})$ for each $j \in [\ell]$. Observe that by construction, $p^{\lambda}$ is the smallest value such that $p^{\lambda} \cdot p^{\beta_j + \sum_{x_i \in vars(S_j) - \mathbf{x}_j} e_i^{(\mathcal{B})}} \geq m_j$ for all $j$ and $\sum_{i \in \mathbf{x}} e_i^{(\mathcal{B})} \leq 1 - \alpha$. Lemma 4.3 then implies that the load due to relation $S_j^{(\mathcal{B})}(\mathbf{h})$ is at most a polylogarithmic factor times

$$\max(m_j/p, m_j / p^{\beta_j + \sum_{x_i \in vars(S_j) - H_j} e_i^{(\mathcal{B})}})$$

which is at most $\max(m_j/p, p^{\lambda})$. □

**Lemma 4.5.** *Every tuple in the join of* $(S_j)_{j \in [\ell]}$ *is contained in a join of subrelations* $(S_j^{(\mathcal{B})})_{j \in [\ell]}$ *for some bin combination* $\mathcal{B}$*.*

*Proof.* Observe first that every join tuple is consistent with the empty bin combination $\mathcal{B}_\emptyset$. Therefore the join of $(S_j^{(\mathcal{B}_\emptyset)})_{j\in[\ell]}$ contains all join tuples that do not contain an overweight heavy hitter $\mathbf{h}_j$ for any relation $S_j$ with respect to $\mathcal{B}_\emptyset$ (and therefore contains all join tuples that are not consistent with any heavy hitter). Now fix a join tuple $\mathbf{t}$ that is overweight for $\mathcal{B}_\emptyset$. By definition, there is an associated relation $S_{j_1}$ and $\mathbf{x}^1 \subset vars(S_{j_1})$ such that $\mathbf{h}^1 = (\mathbf{t}_{\mathbf{x}^1})$ is an overweight heavy hitter of $S_{j_1}$ for $\mathcal{B}_\emptyset$. Let $\mathcal{B}_1$ be the bin combination associated with $\mathbf{h}^1$. By definition $\mathbf{h}^1 \in C'(\mathcal{B}_1)$. Now either $\mathbf{t}$ is contained in the join of $(S_j^{(\mathcal{B}_1)})_{j\in[\ell]}$ and we are done or there is some relation $S_{j_2}$ and $\mathbf{x}^2$ such that $\mathbf{x}^2 - \mathbf{x}^1 \subset vars(S_{j_2})$ such that $\mathbf{h}^2 = (\mathbf{t}_{\mathbf{x}^2})$ has the property that $\mathbf{h}^2_{j_2}$ is an overweight heavy hitter of $S_{j_2}$ for $\mathcal{B}_1$. Again, in the latter case, if $\mathcal{B}_2$ is the bin combination associated with $\mathbf{h}^2$ then $\mathbf{h}^2 \in C'(\mathcal{B}_2)$ by definition and we can repeat the previous argument for $\mathcal{B}_2$ instead of $\mathcal{B}_1$. Since the number of variables grows at each iteration, we can repeat this at most $k$ times before finding a first $\mathcal{B}_r$ such that $\mathbf{t}$ is not associated with any overweight heavy hitter for $\mathcal{B}_r$. In this case $\mathbf{t}$ will be computed in the join of $(S_j^{(\mathcal{B}_r)})_{j\in[\ell]}$. □

We can now prove the main theorem.

**Theorem 4.6.** *The algorithm that computes the joins of $(S_j^{(\mathcal{B})})_{j\in[\ell]}$ for every bin combination $\mathcal{B}$ using the HC algorithm as described above has maximum load $L \leq \log^{O(1)} p \cdot \max_{\mathcal{B}} p^{\lambda(\mathcal{B})}$.*

*Proof.* There are only $\log^{O(1)} p$ choices of $\mathcal{B}$ and for each choice of $\mathcal{B}$, by Corollary 4.4, the load with $p$ virtual processors is $\log^{O(1)} p \cdot \max(L_{min}, p^{\lambda(\mathcal{B})})$. To derive our claim it suffices to show that we can remove the $L_{min}$ term. Observe that in the original LP which corresponds to an empty bin combination $\mathcal{B}$, we have $\lambda + \sum_{i\in S_j} e_i \geq \mu_j$ for each $j \in [\ell]$ and $\sum_i e_i \leq 1$. This implies that $\lambda \geq \mu_j - 1$ and hence $p^\lambda \geq m_j/p$ for each $j$, so $p^\lambda \geq L_{min}$. □

To show the optimality of the algorithm, we apply Theorem 4.7 (which we prove in the next subsection) for a particular bin combination $\mathcal{B}$ with set $\mathbf{x}$. Notice that $m_j(\mathbf{h}_j) \geq m_j/(2p^{\beta_j})$ and further the number of $\mathbf{h}$ is at most $p^\alpha$. Hence the lower bound we obtain from Equation 12 is of the form $L \geq (\prod_j (m_j/p^{\beta_j})^{u_j}/p^{1-\alpha})^{1/u}/2$ for any packing $\mathbf{u}$ of $q_{\mathbf{x}}$ that saturates $\mathbf{x}$, which can be easily seen to be $p^{\lambda(\mathcal{B})}/2$, using a duality argument with the LP in (11).

### 4.3 Lower Bound

In this section we give a lower bound for the load of any deterministic algorithm that computes a query $q$, and generalizes the lower bound in Theorem 3.5, which was over databases with cardinality statistics $\mathbf{M}$. Our new lower bound generalizes this to databases with a fixed degree sequence: if the degree sequence is skewed, then the new bounds can be stronger, proving that skew in the input data makes query evaluation harder.

Let $\mathbf{x}_j \subseteq vars(S_j)$ and $d_j = |\mathbf{x}_j|$. A *statistics of type* $\mathbf{x}_j$ is a function $m_j : [n]^{d_j} \to \mathbb{N}$. An instance $S_j$ satisfies the statistics $m_j$ if for any tuple $\mathbf{h}_j \in [n]^{d_j}$, its frequency is precisely $m_j(\mathbf{h}_j)$, in other words $|\sigma_{\mathbf{x}_j=\mathbf{h}_j}(S_j)| = m_j(\mathbf{h}_j)$. For an example, if $S(x, y)$ is a binary relation, then an $x$-statistics is a degree sequence $m(1), m(2), \ldots, m(n)$; for another example, if $\mathbf{x}_j = \emptyset$ then an $\mathbf{x}$-statistics consists of a single number, which denotes the cardinality of $S_j$. In general, the $\mathbf{x}_j$ statistics define uniquely the cardinality of $S_j$, as $|S_j| = \sum_{\mathbf{h}_j\in[n]^{d_j}} m_j(\mathbf{h}_j)$.

Fix a set of variables $\mathbf{x}$ from $q$, let $d = |\mathbf{x}|$, and denote $\mathbf{x}_j = \mathbf{x} \cap vars(S_j)$ for every $j$. A *statistics of type* $\mathbf{x}$ *for the database* is a vector $\mathbf{m} = (m_1, \ldots, m_\ell)$, where each $m_j$ is an $\mathbf{x}_j$-statistics for $S_j$. We associate with $m$ the function $m : [n]^k \to (\mathbb{N})^\ell$, $m(\mathbf{h}) = (m_1(\mathbf{h}_1), \ldots, m_\ell(\mathbf{h}_\ell))$; here and in the rest of the section, $\mathbf{h}_j$ denotes $\pi_{\mathbf{x}_j}(\mathbf{h})$, *i.e.* the restriction of $\mathbf{h}$ to the variables in $\mathbf{x}_j$. When $\mathbf{x} = \emptyset$, then $m$ simply consists of $\ell$ numbers, each representing the cardinality of a relation; thus, a $\mathbf{x}$-statistics generalizes the cardinality statistics from section 3. Recall that we use upper case $\mathbf{M} = (M_1, \ldots, M_\ell)$ to denote the same statistics expressed in bits, i.e. $M_j(\mathbf{h}_j) = a_j m_j(\mathbf{h}_j) \log n$. As before, we fix some constant $0 < \delta < 1$, and assume every relation $S_j$, has cardinality $\leq n^\delta$.

In this section, we fix statistics $\mathbf{M}$ of type $\mathbf{x}$ and consider the probability space where the instance is chosen uniformly at random over all instances that satisfy $\mathbf{M}$.

To prove the lower bound we need some notations. Let $q_{\mathbf{x}}$ be the *residual query*, obtained by removing all variables $\mathbf{x}$, and decreasing the arities of $S_j$ as necessary (the new arity of $S_j$ is $a_j - d_j$). Clearly, every fractional edge packing of $q$ is also a fractional edge packing of $q_{\mathbf{x}}$, but the converse does not hold in general. Let $\mathbf{u}$ be a fractional edge packing of $q_{\mathbf{x}}$. We say that $\mathbf{u}$ *saturates* a variable $x_i \in \mathbf{x}$, if $\sum_{j:x_i \in vars(S_j)} u_j \geq 1$; we say that $\mathbf{u}$ saturates $\mathbf{x}$ if it saturates all variables in $\mathbf{x}$. For every fractional edge packing $\mathbf{u}$ of $q_{\mathbf{x}}$ that saturates $\mathbf{x}$, denote $u = \sum_{j=1}^{\ell} u_j$ and, using $K$ defined in (6):

$$L_{\mathbf{x}}(\mathbf{u}, \mathbf{M}, p) = \left( \frac{\sum_{\mathbf{h} \in [n]^d} K(\mathbf{u}, \mathbf{M}(\mathbf{h}))}{p} \right)^{1/u} \tag{12}$$

Further denote $L_{\texttt{lower}} = \max_{\mathbf{u}} L_{\mathbf{x}}(\mathbf{u}, \mathbf{M}, p)$, and let $c$ be the constant $c = \min_j \frac{a_j - d_j - \delta}{3a_j}$.

**Theorem 4.7.** *Fix statistics* $\mathbf{M}$ *of type* $\mathbf{x}$*, and consider any deterministic MPC algorithm that runs in one communication round on* $p$ *servers and has maximum load* $L$*. Then, for any edge packing* $\mathbf{u}$ *of* $q$ *that saturates* $\mathbf{x}$*, any algorithm that computes* $q$ *correctly must have load* $L \geq cL_{\mathbf{x}}(\mathbf{u}, \mathbf{M}, p)$ *bits.*

Note that, when $\mathbf{x} = \emptyset$ then $L_{\mathbf{x}}(\mathbf{u}, \mathbf{M}, p) = L(\mathbf{u}, \mathbf{M}, p)$, defined in (7); therefore, our theorem is a generalization of the simpler lower bound Theorem 3.5. Before we prove the theorem, we show an example.

**Example 4.8.** *Consider* $q(x, y, z) = S_1(x, z), S_2(y, z)$*. If* $\mathbf{x} = \emptyset$ *then the lower bound is* $\max_{\mathbf{u}} L(\mathbf{u}, (M_1, M_2), p)$*, which is* $\max(M_1/p, M_2/p)$*, because there are two packings in* $pk(q)$*:* $(1, 0)$ *and* $(0, 1)$*. For* $\mathbf{x} = \{z\}$*, the residual query is* $q_{\mathbf{x}} = S_1(x), S_2(y)$ *and its sole packing is* $(1, 1)$*, which saturates the variable* $z$*. This gives us a new lower bound:* $\sqrt{\sum_{h \in [n]} M_1(h) M_2(h)/p}$*. The lower bound is the maximum of these quantities, and we have given in section 4 an algorithm that matches this bound.*

*Next, consider* $C_3$*. In addition to the lower bounds in Example 3.7, we obtain new bounds by setting* $\mathbf{x} = \{x_1\}$*. The residual query is* $S_1(x_2), S_2(x_2, x_3), S_3(x_3)$ *and* $(1, 0, 1)$ *is a packing that saturates* $x_1$ *(while for example* $(0, 1, 0)$ *does not). This gives us a new lower bound, of the form* $\sqrt{\sum_{h \in [n]} m_1(h) m_3(h)/p}$*.*

In the rest of the section we prove Theorem 4.7.

Let $J(\mathbf{x}) = \{j \mid \mathbf{x} \cap vars(S_j) \neq \emptyset\}$ denote the set of relations that contain some variable from $\mathbf{x}$. Fix a concrete instance $S_j$, and let $\mathbf{a}_j \in S_j$. We write $\mathbf{a}_j | \mathbf{h}$ to denote that the tuple $\mathbf{a}_j$ from $S_j$ matches with $\mathbf{h}$ at their common variables, and denote $(S_j)_{\mathbf{h}}$ the subset of tuples $\mathbf{a}_j$ that match $\mathbf{h}$: $(S_j)_{\mathbf{h}} = \{\mathbf{a}_j \mid \mathbf{a}_j \in S_j, \mathbf{a}_j | \mathbf{h}\}$. Let $I_{\mathbf{h}}$ denote the restriction of $I$ to $\mathbf{h}$, in other words $I_{\mathbf{h}} = ((S_1)_{\mathbf{h}}, \ldots, (S_\ell)_{\mathbf{h}})$.

When $I$ is chosen at random over the probability space defined by the $\mathbf{x}$-statistics $\mathbf{M}$, then, for a fixed tuple $\mathbf{h} \in [n]^d$, the restriction $I_{\mathbf{h}}$ is a uniformly chosen instance over all instances with cardinalities $\mathbf{M}(\mathbf{h})$, which is precisely the probability space that we used in the proof of Theorem 3.6. In particular, for every $\mathbf{a}_j \in [n]^{d_j}$, the probability that $S_j$ contains $\mathbf{a}_j$ is $P(\mathbf{a}_j \in S_j) = m_j(\mathbf{h}_j)/n^{a_j-d_j}$; thus, our proof below is an extension of that of Theorem 3.6.

Lemma A.1 immediately gives:

**Lemma 4.9.** $\mathbf{E}[|q(I_{\mathbf{h}})|] = n^{k-d}\prod_{j=1}^{\ell}\frac{m_j(\mathbf{h}_j)}{n^{a_j-d_j}}$

Let us fix some server and let $m(I)$ be the message the server receives on input $I$. For any fixed value $m$ of $m(I)$, let $K_m(S_j)$ be the set of tuples from relation $S_j$ *known* by the server. Let $w_j(\mathbf{a}_j)$ to denote the probability that the server knows the tuple $\mathbf{a}_j \in S_j$. In other words $w_j(\mathbf{a}_j) = P(\mathbf{a}_j \in K_{m_j(S_j)}(S_j))$, where the probability is over the random choices of $S_j$. This is obviously upper bounded by $P(\mathbf{a}_j \in S_j)$:

$$w_j(\mathbf{a}_j|\mathbf{h}) \leq m_j(\mathbf{h}_j)/n^{a_j-d_j} \tag{13}$$

We derive a second upper bound by exploiting the fact that the server receives a limited number of bits:

**Lemma 4.10.** *Let $L$ be the number of bits a server receives. Then, $\sum_{\mathbf{a}_j\in[n]^{a_j}} w_j(\mathbf{a}_j) \leq \frac{L}{ca_j\log(n)}$.*

*Proof.* Since $\sum_{\mathbf{a}_j\in[n]^{a_j}} w_j(\mathbf{a}_j) = \mathbf{E}[|K_{m(S_j)}(S_j)|]$, we will bound the right hand side. Now, notice that:

$$\begin{aligned} H(S_j) &= H(m(S_j)) + \sum_m P(m(S_j)=m)\cdot H(S_j|m(S_j)=m) \\ &\leq L + \sum_m P(m(S_j)=m)\cdot H(S_j|m(S_j)=m) \end{aligned} \tag{14}$$

For every $\mathbf{h}$, let $K_m((S_j)_{\mathbf{h}})$ denote the known tuples from the restriction of $S_j$ to $\mathbf{h}$. Following the proof of Lemma A.2, for the constant $c = \log e + 1$,

$$\begin{aligned} H(S_j|m(S_j)=m) &\leq \sum_{\mathbf{h}}\left(1-\frac{|K_m((S_j)_{\mathbf{h}})|}{cm_j(\mathbf{h}_j)}\right)\log\binom{n^{a_j-d_j}}{m_j(\mathbf{h}_j)} \\ &= H(S_j) - \sum_{\mathbf{h}}\frac{|K_m((S_j)_{\mathbf{h}})|}{cm_j(\mathbf{h}_j)}\log\binom{n^{a_j-d_j}}{m_j(\mathbf{h}_j)} \\ &\leq H(S_j) - \sum_{\mathbf{h}}\frac{|K_m((S_j)_{\mathbf{h}})|}{cm_j(\mathbf{h}_j)}m_j(\mathbf{h}_j)(a_j-d_j-\delta)\log(n) \\ &= H(S_j) - (1/c)\cdot|K_m(S_j)|(a_j-d_j-\delta)\log(n) \end{aligned}$$

Plugging this in Equation 14, we have:

$$H(S_j) \leq L + H(S_j) - (1/c)\cdot\mathbf{E}[|K_m(S_j)|](a_j-d_j-\delta)\log(n)$$

or equivalently, since $c \leq 3$:

$$\mathbf{E}[|K_m(S_j)|] \leq \frac{3L}{(a_j-d_j-\delta)\log(n)}$$

This concludes our proof. □

Let $q_{\mathbf{x}}$ be the residual query, and recall that $\mathbf{u}$ is a fractional edge packing that saturates $\mathbf{x}$. Define the *extended query* $q_{\mathbf{x}}'$ to consists of $q_{\mathbf{x}}$, where we add a new atom $S_i'(x_i)$ for every variable $x_i \in vars(q_{\mathbf{x}})$. Define $u_i' = 1 - \sum_{j:i\in S_j} u_j$. In other words, $u_i'$ is defined to be the slack at the variable $x_i$ of the packing $\mathbf{u}$. The new edge packing $(\mathbf{u}, \mathbf{u}')$ for the extended query $q_{\mathbf{x}}'$ has no more slack, hence it is both a tight fractional edge packing and a tight fractional edge cover for $q_{\mathbf{x}}$. By adding all equalities of the tight packing we obtain:

$$\sum_{j=1}^{\ell}(a_j - d_j)u_j + \sum_{i=1}^{k-d} u_i' = k - d$$

We next compute how many output tuples from $q(I_{\mathbf{h}})$ will be known in expectation by the server. Note that $q(I_{\mathbf{h}}) = q_{\mathbf{x}}(I_{\mathbf{h}})$, and thus:

$$\begin{aligned}\mathbf{E}[|K_m(q(I_{\mathbf{h}}))|] = \mathbf{E}[|K_m(q_{\mathbf{x}}(I_{\mathbf{h}}))|] &= \sum_{\mathbf{a}\in[n]^{k-d}}\prod_{j=1}^{\ell} w_j(\mathbf{a}_j|\mathbf{h}) \\ &= \sum_{\mathbf{a}\in[n]^{k-d}}\prod_{j=1}^{\ell} w_j(\mathbf{a}_j|\mathbf{h})\prod_{i=1}^{k-d} w_i'(\mathbf{a}_i) \\ &\le \prod_{i=1}^{k-d} n^{u_i'}\cdot\prod_{j=1}^{\ell}\left(\sum_{\mathbf{a}\in[n]^{a_j-d_j}} w_j(\mathbf{a}|\mathbf{h})^{1/u_j}\right)^{u_j}\end{aligned}$$

By writing $w_j(\mathbf{a}|\mathbf{h})^{1/u_j} = w_j(\mathbf{a}|\mathbf{h})^{1/u_j-1}w_j(\mathbf{a}|\mathbf{h})$, we can bound the quantity as follows:

$$\begin{aligned}\sum_{\mathbf{a}\in[n]^{a_j-d_j}} w_j(\mathbf{a}|\mathbf{h})^{1/u_j} &\le \left(\frac{m_j(\mathbf{h}_j)}{n^{a_j-d_j}}\right)^{1/u_j-1}\sum_{\mathbf{a}\in[n]^{a_j-d_j}} w_j(\mathbf{a}|\mathbf{h}) \\ &= (m_j(\mathbf{h}_j)n^{d_j-a_j})^{1/u_j-1}L_j(\mathbf{h})\end{aligned}$$

where $L_j(\mathbf{h}) = \sum_{\mathbf{a}\in[n]^{a_j-d_j}} w_j(\mathbf{a}|\mathbf{h})$. Notice that for every relation $S_j$, $\sum_{\mathbf{h}_j\in[n]^{d_j}} L_j(\mathbf{h}_j) = \sum_{\mathbf{a}_j\in[n]^{a_j}} w_j(\mathbf{a}_j)$.

$$\begin{aligned}\mathbf{E}[|K_m(q(I_{\mathbf{h}}))|] &\le n^{\sum_{i=1}^{k-d}u_i'}\prod_{j=1}^{\ell}\left(L_j(\mathbf{h})m_j(\mathbf{h}_j)^{1/u_j-1}n^{(d_j-a_j)(1/u_j-1)}\right)^{u_j} \\ &= \prod_{j=1}^{\ell}L_j(\mathbf{h})^{u_j}\cdot\prod_{j=1}^{\ell}m_j(\mathbf{h}_j)^{-u_j}\cdot\mathbf{E}[|q(I_{\mathbf{h}})|] \qquad (15)\end{aligned}$$

Summing over all $p$ servers, we obtain that the expected number of answers that can be output for $q(I_{\mathbf{h}})$ is at most $p\mathbf{E}[|K_m(q(I_{\mathbf{h}}))|]$. If some $\mathbf{h} \in [n]^d$ this number is not at least $\mathbf{E}[|q(I_{\mathbf{h}})|]$, the algorithm will fail to compute $q(I)$. Consequently, for every $\mathbf{h}$ we must have that $\prod_{j=1}^{\ell} L_j(\mathbf{h}_j)^{u_j} \ge (1/p)\cdot\prod_{j=1}^{\ell} m_j(\mathbf{h}_j)^{u_j}$. Summing the inequalities for every $\mathbf{h} \in [n]^d$:

$$(1/p)\cdot\sum_{\mathbf{h}}\prod_{j=1}^{\ell} m_j(\mathbf{h}_j)^{u_j} \le \sum_{\mathbf{h}}\prod_{j=1}^{\ell} L_j(\mathbf{h}_j)^{u_j}$$

To finish the lower bound, we will apply Friedgut's inequality on the right hand side of the above inequality.

**Lemma 4.11.** *If* $\mathbf{u}$ *saturates* $\mathbf{x}$ *in query* $q$*,*

$$\sum_{\mathbf{h}} \prod_{j=1}^{\ell} L_j(\mathbf{h}_j)^{u_j} \leq \prod_{j=1}^{\ell} \left( \sum_{\mathbf{h}_j} L_j(\mathbf{h}_j) \right)^{u_j}$$

Since we have $M_j = a_j m_j \log(n)$, and $\sum_{\mathbf{h}_j} L_j(\mathbf{h}_j) \leq L/(ca_j \log(n))$ (from Lemma 4.10), we obtain that for any edge packing $\mathbf{u}$ that saturates $\mathbf{x}$, it must be that

$$L \geq c \cdot \left( \frac{\sum_{\mathbf{h} \in [n]^d} \prod_j M_j(\mathbf{h}_j)^{u_j}}{p} \right)^{1/u}$$

# 5 Map-Reduce Models

In this section, we discuss the connection between the MPC model and the Map-Reduce model presented by Afrati et. al [1]. In contrast to the MPC model, where the number of servers $p$ is the main parameter, in the model of [1] the main parameter is an upper bound $q$ on the number of input tuples a reducer can receive, which is called *reducer size*. Given an input $I$, a Map-Reduce algorithm is restricted to deterministically send each input tuple independently to some reducer, which will then produce all the outputs that can be extracted from the received tuples. If $q_i \leq q$ is the number of inputs assigned to the $i$-th reducer, where $i = 1, \ldots, p$, we define the *replication rate* $r$ of the algorithm $r = \sum_{i=1}^{p} q_i/|I|$.

In [1], the authors provide lower and upper bounds on $r$ with respect to $q$ and the size of the input. However, their results are restricted to binary relations where all sizes are equal, and they provide matching upper and lower bounds only for a subclass of such queries. We show next how to apply the results of this paper to remove these restrictions, and further consider an even stronger computation model for our lower bounds.

First, we express the bound on the data received by the reducers in *bits*: let $L$ be the maximum number of bits each reducer can receive. The input size $|I|$ is also expressed in bits. We will also allow any algorithm to use randomization. Finally, we relax the assumption on how the inputs are communicated to the reducers: instead of restricting each input tuple to be sent independently, we assume that the input data is initially partitioned into a number of *input servers* $p_0$ (where $p_0$ must be bigger than the query size), and allow the algorithm to communicate bits to reducers by accessing all the data in one such input server. Notice that this setting allows for stronger algorithms that use input statistics to improve communication.

If each reducer receives $L_i$ bits, the replication rate $r$ is defined as $r = \sum_{i=1}^{p} L_i/|I|$. Notice that any algorithm with replication rate $r$ must use $p \geq (r|I|)/L$ reducers. Now, let $q$ be a conjunctive query, where $S_j$ has $M_j = a_j m_j \log(n)$ bits ($n$ is the size of the domain).

**Theorem 5.1.** *Let $q$ be a conjunctive query where $S_j$ has size (in bits) $M_j$. Any algorithm that computes $q$ with reducer size $L$, where $L \leq M_j$ for every $S_j$[5] must have replication rate*

$$r \geq \frac{c^u L}{\sum_j M_j} \max_{\mathbf{u}} \prod_{j=1}^{\ell} \left(\frac{M_j}{L}\right)^{u_j}$$

*where* $\mathbf{u}$ *ranges over all fractional edge packings of $q$.*

*Proof.* Let $f_i$ be the fraction of answers returned by server $i$, in expectation, where $I$ is a randomly chosen database with statistics $\mathbf{M}$. By Theorem 3.5, $f_i \leq \frac{L_i^u}{c^u K(\mathbf{u},\mathbf{M})}$. Since we assume all answers are returned,

$$1 \leq \sum_{i=1}^{p} f_i = \sum_{i=1}^{p} \frac{L_i^u}{c^u K(\mathbf{u},\mathbf{M})} = \frac{\sum_{i=1}^{p} L_i L_i^{u-1}}{c^u K(\mathbf{u},\mathbf{M})} \leq \frac{L^{u-1}\sum_{i=1}^{p} L_i}{c^u K(\mathbf{u},\mathbf{M})} = \frac{L^{u-1} r|I|}{c^u K(\mathbf{u},\mathbf{M})}$$

where we used the fact that $u \geq 1$ for the optimal $\mathbf{u}$. The claim follows by using the definition of $K$ (6) and noting that $|I| = \sum_j M_j$. □

We should note here that, following from our analysis in section 3, the bound provided by Theorem 5.1 is matched by the HYPERCUBE algorithm with appropriate shares. We illustrate Theorem 5.1 with an example.

**Example 5.2** (Triangles)**.** *Consider the triangle query $C_3$ and assume that all sizes are equal to $M$. In this case, the edge packing that maximizes the lower bound is the one that maximizes $\sum_j u_j$, $(1/2, 1/2, 1/2)$. Thus, we obtain a bound $\Omega(\sqrt{M/L})$ for the replication rate. This is exactly the formula proved in [1], but notice that we can derive bounds even if the sizes are not equal. Further, observe that any algorithm must use at least $\Omega((M/L)^{3/2})$ reducers.*

## 6 Conclusions

In this paper we have studied the parallel query evaluation problem on databases with known cardinalities, and with known cardianlities and heavy hitters. In the first case we have given matching lower and upper bounds that are described in terms of *fractional edge packings* of the query. In the second case we have also described matching lower and upper bounds described in terms of fractional packings of the residual queries, one residual query for each type of heavy hitter.

## References


[1] F. N. Afrati, A. D. Sarma, S. Salihoglu, and J. D. Ullman. Upper and lower bounds on the cost of a map-reduce computation. *PVLDB*, 6(4):277–288, 2013.

[2] F. N. Afrati and J. D. Ullman. Optimizing joins in a map-reduce environment. In *EDBT*, pages 99–110, 2010.


[5] if $L > M_j$, we can send the whole relation to any reducer without cost

[3] A. Atserias, M. Grohe, and D. Marx. Size bounds and query plans for relational joins. In *FOCS*, pages 739–748, 2008.

[4] P. Beame, P. Koutris, and D. Suciu. Communication steps for parallel query processing. In *PODS*, pages 273–284, 2013.

[5] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In *OSDI*, pages 137–150, 2004.

[6] E. Friedgut. Hypergraphs, entropy, and inequalities. *American Mathematical Monthly*, pages 749–760, 2004.

[7] S. Ganguly, A. Silberschatz, and S. Tsur. Parallel bottom-up processing of datalog queries. *J. Log. Program.*, 14(1&2):101–126, 1992.

[8] M. Grohe and D. Marx. Constraint solving via fractional edge covers. In *SODA*, pages 289–298, 2006.

[9] H. Q. Ngo, E. Porat, C. Ré, and A. Rudra. Worst-case optimal join algorithms: [extended abstract]. In *PODS*, pages 37–48, 2012.

[10] C. Olston, B. Reed, U. Srivastava, R. Kumar, and A. Tomkins. Pig latin: a not-so-foreign language for data processing. In *SIGMOD Conference*, pages 1099–1110, 2008.

[11] S. Suri and S. Vassilvitskii. Counting triangles and the curse of the last reducer. In *WWW*, pages 607–614, 2011.

[12] C. B. Walton, A. G. Dale, and R. M. Jenevein. A taxonomy and performance model of data skew effects in parallel joins. In *VLDB*, pages 537–548, 1991.

[13] R. S. Xin, J. Rosen, M. Zaharia, M. J. Franklin, S. Shenker, and I. Stoica. Shark: Sql and rich analytics at scale. In *SIGMOD Conference*, pages 13–24, 2013.

[14] Y. Xu, P. Kostamaa, X. Zhou, and L. Chen. Handling data skew in parallel joins in shared-nothing systems. In *SIGMOD Conference*, pages 1043–1052, 2008.

[15] A. C. Yao. Lower bounds by probabilistic arguments. In *FOCS*, pages 420–428, Tucson, AZ, 1983.

[16] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica. Resilient distributed datasets: a fault-tolerant abstraction for in-memory cluster computing. In *NSDI*, 2012.

# A Proof of The Lower Bound

In this section, we present the detailed proof of Theorem 3.5. Let $\mathbf{E}[|q(I)|]$ denote the expected number of answers to the query $q$.

**Lemma A.1.** *The expected number of answers to $q$ is $\mathbf{E}[|q(I)|] = n^{k-a}\prod_{j=1}^{\ell} m_j$.*

*Proof.* For any relation $S_j$, and any tuple $\mathbf{a}_j \in [n]^{a_j}$, the probability that $S_j$ contains $\mathbf{a}_j$ is $P(\mathbf{a_j} \in S_j) = m_j/n^{a_j}$. Given a tuple $\mathbf{a} \in [n]^k$ of the same arity as the query answer, let $\mathbf{a}_j$ denote its projection on the variables in $S_j$. Then:

$$\mathbf{E}[|q(I)|] = \sum_{\mathbf{a}\in[n]^k} P(\bigwedge_{j=1}^{\ell}(\mathbf{a}_j \in S_j)) = \sum_{\mathbf{a}\in[n]^k}\prod_{j=1}^{\ell} P(\mathbf{a}_j \in S_j)$$
$$= \sum_{\mathbf{a}\in[n]^k}\prod_{j=1}^{\ell} m_j n^{-a_j} = n^{k-a}\prod_{j=1}^{\ell} m_j$$

□

Let us fix some server and let $m(I)$ denote the function specifying the message the server receives on input $I$. For any fixed value $m$ of $m(I)$, we define the set of tuples of a relation $R$ of arity *a known* by the server given message $m$ as

$$K_m(R) = \{t \in [n]^a \mid \text{ for all } I, m(I) = m \Rightarrow t \in R(I)\}$$

We will particularly apply this definition with $R = S_j$ and $R = q$. Clearly, an output tuple $\mathbf{a} \in K_m(q)$ iff for every $j$, $\mathbf{a}_j \in K_m(S_j)$, where $\mathbf{a}_j$ denotes the projection of $\mathbf{a}$ on the variables in the atom $S_j$.

**Bounding the Knowledge of Each Relation** Let us fix a server, and an input relation $S_j$. Observe that, for a randomly chosen database $I$, $S_j$ is a randomly chosen subset of $[n]^{a_j}$ of size $m_j$. Since there are $\binom{n^{a_j}}{m_j}$ such subsets, the number of bits necessary to represent the relation is $\log\binom{n^{a_j}}{m_j}$.

Let $m(S_j)$ be the part of the message $m$ received from the server that corresponds to $S_j$. The following lemma provides a bound on the expected knowledge $K_{m(S_j)}(S_j)$ the server may obtain from $S_j$:

**Lemma A.2.** *Suppose that for all $\binom{n^{a_j}}{m_j}$ relations $S_j$, where $m_j \le n^{a_j}/2$, $m(S_j)$ is at most $f_j \cdot \log\binom{n^{a_j}}{m_j}$ bits long. Then $\mathbf{E}[|K_{m(S_j)}(S_j)|] \le (\log e + 1) f_j \cdot m_j$.*

*Proof.* Let $m$ be a possible value for $m(S_j)$. Since $m$ fixes precisely $|K_m(S_j)|$ tuples of $S_j$,

$$\log|\{S_j \mid m(S_j) = m\}| \le \log\binom{n^{a_j} - |K_m(S_j)|}{m_j - |K_m(S_j)|}$$

We will apply the following lemma to the above equation.

**Lemma A.3.** *For any $k \leq m \leq N/2$, and constant $c = \log e + 1$,*

$$\log \binom{N-k}{m-k} \leq \left(1 - \frac{k}{cm}\right) \log \binom{N}{m}$$

*Proof.* We have:

$$\log \binom{N-k}{m-k} = \log \binom{N}{m} - \log \frac{[N]_k}{[m]_k}$$

where $[N]_k = N \cdot (N-1) \cdots (N-k+1)$. Since we also have that $\binom{N}{m} = [N]_m / [m]_m$, it suffices to show that $\log([N]_k/[m]_k) \geq k/(cm) \log \binom{N}{m}$, or equivalently $\frac{[N]_k}{[m]_k} \geq \binom{N}{m}^{k/(cm)}$. To show this, notice that the following two inequalities hold:

$$\frac{[N]_k}{[m]_k} \geq (N/m)^k, \quad \binom{N}{m} \leq (eN/m)^m$$

We are left to prove $(N/m)^k \geq (eN/m)^{k/c}$, or equivalently $N/m \geq (eN/m)^{1/c}$, equivalently $(N/m)^{c-1} \geq e$. However, since $m \leq N/2$, we have $N/m \geq 2$ and consequently $(N/m)^{c-1} \geq 2^{c-1} = e$, which proves the lemma. □

We now have, for $c = \log e + 1$:

$$\log |\{S_j \mid m(S_j) = m\}| \leq \left(1 - \frac{|K_m(S_j)|}{cm_j}\right) \log \binom{n^{a_j}}{m_j} \tag{16}$$

We can bound the value we want by considering the binary entropy of the distribution $S_j$. By applying the chain rule for entropy, we have

$$\begin{aligned}
&H(S_j) \\
&= H(m(S_j)) + \sum_m P(m(S_j) = m) \cdot H(S_j | m(S_j) = m) \\
&\leq f_j \cdot H(S_j) + \sum_m P(m(S_j) = m) \cdot H(S_j | m(S_j) = m) \\
&\leq f_j \cdot H(S_j) + \sum_m P(m(S_j) = m) \cdot (1 - \frac{|K_m(S_j)|}{c \cdot m_j}) H(S_j) \\
&= f_j \cdot H(S_j) + (1 - \sum_m P(m(S_j) = m) \, \frac{|K_m(S_j)|}{c \cdot m_j}) H(S_j) \\
&= f_j \cdot H(S_j) + (1 - \frac{\mathbf{E}[|K_{m(S_j)}(S_j)|]}{c \cdot m_j}) H(S_j)
\end{aligned} \tag{17}$$

where the first inequality follows from the assumed upper bound on $|m(S_j)|$, the second inequality follows by (16), and the last two lines follow by definition. Dividing both sides of (17) by $H(S_j)$ since $H(S_j)$ is not zero and rearranging we obtain the required statement. □

**Bounding the Knowledge of the Query** We know that $\log \binom{n^{a_j}}{m_j}$ is the number of bits needed to represent relation $S_j$. Since we want to show a lower bound of $L$, we assume that each server receives at most $L$ bits. Let us also fix some server. The message $m = m(I)$ received by the server is the concatenation of $\ell$ messages, one for each input relation. $K_m(S_j)$ depends only on $m(S_j)$, so we can assume w.l.o.g. that $K_m(S_j) = K_{m(S_j)}(S_j)$.

The message $m(I)$ will contain at most $L$ bits. Now, for each relation $S_j$, let us define

$$f_j = \frac{\max_{S_j} |m(S_j)|}{\log \binom{n^{a_j}}{m_j}}.$$

Thus, $f_j$ is the largest fraction of bits of $S_j$ that the server receives, over all choices of the matching $S_j$. We now derive an upper bound on the $f_j$'s. As we discussed before, each part $m(S_j)$ of the message is constructed independently of the other relations. Hence, it must be that $\sum_{j=1}^{\ell} \max_{S_j} |m(S_j)| \leq L$. By substituting the definition of $f_j$ in this equation, we obtain that

$$L \geq \sum_{j=1}^{\ell} f_j \log \binom{n^{a_j}}{m_j}.$$

Using our assumptions on the choice of $n$, we can write:

$$\log \binom{n^{a_j}}{m_j} \geq m_j \log(n^{a_j}/m_j) \geq m_j(a_j - \delta) \log n$$

Denoting $M_j = a_j m_j \log n$, we now have

$$L \geq \sum_{j=1}^{\ell} \frac{a_j - \delta}{a_j} f_j M_j \geq \min_j \frac{a_j - \delta}{a_j} \sum_j f_j M_j$$

where the quantity $C_0 = \min_j \frac{a_j - \delta}{a_j}$ is a positive constant (since $\delta \leq a_j$).

For $\mathbf{a}_j \in [n]^{a_j}$, let $w_j(\mathbf{a}_j)$ denote the probability that the server knows the tuple $\mathbf{a}_j$. In other words $w_j(\mathbf{a}_j) = P(\mathbf{a}_j \in K_{m_j(S_j)}(S_j))$, where the probability is over the random choices of $S_j$.

**Lemma A.4.** *For any relation $S_j$:*

*(a)* $\forall \mathbf{a}_j \in [n]^{a_j} : w_j(\mathbf{a}_j) \leq m_j/n^{a_j}$, *and*

*(b)* $\sum_{\mathbf{a}_j \in [n]^{a_j}} w_j(\mathbf{a}_j) \leq 3 f_j \cdot m_j$.

Since the server receives a separate message for each relation $S_j$, from a distinct input server, the events $\mathbf{a}_1 \in K_{m_1}(S_1), \ldots, \mathbf{a}_\ell \in K_{m_\ell}(S_\ell)$ are independent, hence:

$$\mathbf{E}[|K_{m(I)}(q)|] = \sum_{\mathbf{a} \in [n]^k} P(\mathbf{a} \in K_{m(I)}(q)) = \sum_{\mathbf{a} \in [n]^k} \prod_{j=1}^{\ell} w_j(\mathbf{a}_j)$$

We now use Friedgut's inequality. Recall that in order to apply the inequality, we need to find a fractional edge cover. Let us pick any fractional edge packing $\mathbf{u} = (u_1, \ldots, u_\ell)$. Given $q$, consider the *extended query*, which has a new unary atom for each variable $x_i$:

$$q'(x_1, \ldots, x_k) = S_1(\bar{x}_1), \ldots, S_\ell(\bar{x}_\ell), T_1(x_1), \ldots, T_k(x_k)$$

For each new symbol $T_i$, define $u'_i = 1 - \sum_{j:x_i \in \text{vars}(S_j)} u_j$. Since $\mathbf{u}$ is a packing, $u'_i \geq 0$. Let us define $\mathbf{u}' = (u'_1, \ldots, u'_k)$.

**Lemma A.5.** *(a) The assignment $(\mathbf{u}, \mathbf{u}')$ is both a tight fractional edge packing and a tight fractional edge cover for $q'$. (b) $\sum_{j=1}^{\ell} a_j u_j + \sum_{i=1}^{k} u'_i = k$*

*Proof.* (a) is straightforward, since for every variable $x_i$ we have $u'_i + \sum_{j:x_i \in \text{vars}(S_j)} u_j = 1$. Summing up:

$$k = \sum_{i=1}^{k} \left( u'_i + \sum_{j:x_i \in \text{vars}(S_j)} u_j \right) = \sum_{i=1}^{k} u'_i + \sum_{j=1}^{\ell} a_j u_j$$

which proves (b). □

We will apply Friedgut's inequality to the extended query. Set the variables $w(-)$ used in Friedgut's inequality as follows:

$$\begin{aligned} w_j(\mathbf{a}_j) &= P(\mathbf{a}_j \in K_{m_j(S_j)}(S_j)) \text{ for } S_j, \text{ tuple } \mathbf{a}_j \in [n]^{a_j} \\ w'_i(\alpha) &= 1 \qquad\qquad\qquad\qquad \text{for } T_i, \text{ value } \alpha \in [n] \end{aligned}$$

Recall that, for a tuple $\mathbf{a} \in [n]^k$ we use $\mathbf{a}_j \in [n]^{a_j}$ for its projection on the variables in $S_j$; with some abuse, we write $\mathbf{a}_i \in [n]$ for the projection on the variable $x_i$. Assume first that $u_j > 0$, for $j = 1, \ell$. Then:

$$\begin{aligned} \mathbf{E}[|K_m(q)|] &= \sum_{\mathbf{a} \in [n]^k} \prod_{j=1}^{\ell} w_j(\mathbf{a}_j) \\ &= \sum_{\mathbf{a} \in [n]^k} \prod_{j=1}^{\ell} w_j(\mathbf{a}_j) \prod_{i=1}^{k} w'_i(\mathbf{a}_i) \\ &\leq \prod_{j=1}^{\ell} \left( \sum_{\mathbf{a} \in [n]^{a_j}} w_j(\mathbf{a})^{1/u_j} \right)^{u_j} \prod_{i=1}^{k} \left( \sum_{\alpha \in [n]} w'_i(\alpha)^{1/u'_i} \right)^{u'_i} \\ &= \prod_{j=1}^{\ell} \left( \sum_{\mathbf{a} \in [n]^{a_j}} w_j(\mathbf{a})^{1/u_j} \right)^{u_j} \prod_{i=1}^{k} n^{u'_i} \end{aligned}$$

Note that, since $w'_i(\alpha) = 1$ we have $w'_i(\alpha)^{1/u'_i} = 1$ even if $u'_i = 0$. Write $w_j(\mathbf{a})^{1/u_j} = w_j(\mathbf{a})^{1/u_j - 1} w_j(\mathbf{a})$,

and use Lemma A.4 to obtain:

$$\begin{aligned}\sum_{\mathbf{a}\in[n]^{a_j}} w_j(\mathbf{a})^{1/u_j} &\le (m_j/n^{a_j})^{1/u_j-1} \sum_{\mathbf{a}\in[n]^{a_j}} w_j(\mathbf{a}) \\ &\le (m_j n^{-a_j})^{1/u_j-1} 3f_j \cdot m_j \\ &= 3f_j \cdot m_j^{1/u_j} \cdot n^{(a_j - a_j/u_j)}\end{aligned}$$

Plugging this in the bound, we have shown that:

$$\begin{aligned}\mathbf{E}[|K_m(q)|] &\le \prod_{j=1}^{\ell} (3f_j \cdot m_j^{1/u_j} \cdot n^{(a_j-a_j/u_j)})^{u_j} \cdot \prod_{i=1}^{k} n^{u_i'} \\ &= \prod_{j=1}^{\ell} (3f_j)^{u_j} \cdot \prod_{j=1}^{\ell} m_j \cdot n^{(\sum_{j=1}^{\ell} a_j u_j - a)} \cdot n^{\sum_{i=1}^{k} u_i'} \\ &= \prod_{j=1}^{\ell} (3f_j)^{u_j} \cdot \prod_{j=1}^{\ell} m_j \cdot n^{-a+(\sum_{j=1}^{\ell} a_j u_j + \sum_{i=1}^{k} u_i')} \\ &= \prod_{j=1}^{\ell} (3f_j)^{u_j} \cdot \prod_{j=1}^{\ell} m_j \cdot n^{k-a} \\ &= \prod_{j=1}^{\ell} (3f_j)^{u_j} \cdot \mathbf{E}[|q(I)|] \end{aligned} \tag{18}$$

If some $u_j = 0$, then we can derive the same lower bound as follows: We can replace each $u_j$ with $u_j + \delta$ for any $\delta > 0$ still yielding an edge cover. Then we have $\sum_j a_j u_j + \sum_i u_i' = k + a\delta$, and hence an extra factor $n^{a\delta}$ multiplying the term $n^{k-a}$ in (18); however, we obtain the same upper bound since, in the limit as $\delta$ approaches 0, this extra factor approaches 1.

Let $f_q = \prod_{j=1}^{\ell} (3f_j)^{u_j}$; the final step is to upper bound the quantity $f_q$ using the fact that $\sum_{j=1}^{\ell} f_j M_j \le L/C_0$. Indeed:

$$\begin{aligned} f_q = \prod_{j=1}^{\ell} (3f_j)^{u_j} &= \prod_{j=1}^{\ell} \left(\frac{f_j M_j}{u_j}\right)^{u_j} \prod_{j=1}^{\ell} \left(\frac{3u_j}{M_j}\right)^{u_j} \\ &\le \left(\frac{\sum_{j=1}^{\ell} f_j M_j}{\sum_j u_j}\right)^{\sum_j u_j} \prod_{j=1}^{\ell} \left(\frac{3u_j}{M_j}\right)^{u_j} \\ &\le \left(\frac{L/C_0}{\sum_j u_j}\right)^{\sum_j u_j} \prod_{j=1}^{\ell} \left(\frac{3u_j}{M_j}\right)^{u_j} \\ &= \prod_{j=1}^{\ell} \left(\frac{L}{M_j}\right)^{u_j} \prod_{j=1}^{\ell} \left(\frac{3u_j}{C_0 \sum_j u_j}\right)^{u_j} \end{aligned}$$

Here, the first inequality comes from the weighted version of the Arithmetic Mean-Geometric

Mean inequality. Notice also that, since $u_j \leq 1$, we have that $(u_j / \sum_j u_j)^{u_j} \leq 1$, and thus:

$$\prod_{j=1}^{\ell} \left( \frac{3u_j}{C_0 \sum_j u_j} \right)^{u_j} \leq (3/C_0)^{\sum_j u_j}$$

Recall that $u = \sum_j u_j$, therefore,

$$\mathbf{E}[|K_m(q)|] \leq (3/C_0)^u \prod_{j=1}^{\ell} \left( \frac{L}{M_j} \right)^{u_j} = (3/C_0)^u \frac{L^u}{K(\mathbf{u}, \mathbf{M})}$$

which completes the proof of Theorem 3.5. (Recall that our $L$ denotes the load of an arbitrary server, which was denoted $L_i$ in the statement of the theorem.)

# B Hashing

In this section, we discuss hashing and in particular prove Lemma 3.1. Let $R(A_1, \ldots, A_r)$ be a relation of arity $r$ of size $m$. Further, let $[n]$ be the domain for each attribute. For any $i \in [n]$, we define its degree on attribute $A_j$ as $d_j(i) = |\{t \in R \mid t[A_j] = i\}|$. We extend this definition to sets of attributes for $S \subseteq [r]$ and for $I = (i_j)_{j \in S} \in [n]^S$, we let $R_S(I) = \{t \in R \mid t[A_j] = i_j, \forall j \in [S]\}$ and We write $d_S(I) = |R_S(I)|$.

In our setting, we wish to hash each tuple from $R$ on $p = p_1 \times \cdots \times p_r$ bins: each bin corresponds to a point in the $r$-dimensional space $\mathcal{A} = [p_1] \times \ldots [p_r]$. For each attribute $A_j$, we are equipped with a perfect hash function $h_j : [n] \to [p_j]$. Each tuple $t \in R$ will be sent to the bin $B(t)$ as follows:

$$B(t) = (h_1(t[A_1]), \ldots, h_r(t[A_r]))$$

where $t[A_j]$ is the value of the tuple $t$ at attribute $A_j$. For a bin $b \in \mathcal{A}$, let $b_j$ denote its coordinate at the $j$-th dimension. Moreover, let $L(b)$ denote its *load*, in other words:

$$L(b) = |\{t \in R \mid B(t) = b\}|$$

**Lemma B.1.** *For every bin $b \in \mathcal{A}$, $\mathbf{E}[L(b)] = m/p$.*

*Proof.* First, note that for a given tuple $t \in R$, $Pr[B(t) = b] = \prod_{j=1}^{r} Pr[t[A_j] = b_j] = \prod_{j=1}^{r} (1/p_j) = 1/p$. Since we have $m$ tuples, using linearity of expectation we obtain the desired result. □

Thus, the tuples are in expectation distributed equally among the bins. However, this does not tell us anything about the maximum load of the bins, $L = \max_b \{L(b)\}$. An easy observation about $L$ is that $m/p \leq L \leq m$. Ideally, we want $L$ to be equal to the expected load $m/p$, but this is not always possible, as the next example shows.

**Example B.2.** *Let $R = \{(1, \ldots, 1, i) \mid i \in [m]\}$. In this case, every bin that receives a tuple from $R$ will have the same first $r - 1$ coordinates, and hence the $m$ tuples will be distributed among only $p_r$ servers. This implies that $L \geq m/p_r$. Using similar reasoning, there will be some instance of $R$ such that $L \geq m/\min_j\{p_j\}$.*

Let us set up some notation. Let $X_{b,t}$ denote the indicator for the event that the tuple $t$ ends up in bin $b$. Then, $L(b) = \sum_{t \in R} X_{b,t}$. Additionally, for $j = 1, \ldots, r$, and $b_j \in [p_j]$, let $L^j(b_j) = \sum_{b' : b'_j = b_j} L(b')$. Finally, let $L^j = \max_{b_j \in [p_j]} \{L^j(b_j)\}$.

We start our analysis with the case where for every attribute $A_j$, and every $i \in [n]$, the degree $d_j(i)$ is either 0 or 1, or equivalently, each value appears at most once in every attribute. In this case, $R$ is a *matching* relation.

**Lemma B.3.** *Assume that, for every attribute $A$ and every value $i \in [n]$, $d_A(i) \leq 1$. Then,*

$$Pr[L > 3m/p] \leq pe^{-m/p}.$$

*Proof.* Since the tuples contains disjoint values, each tuple is hashed independently and uniformly at random to one of the $p$ bins. Thus, we can apply directly Corollary C.2 to derive the result. □

**Lemma B.4.** *Suppose that for some attribute $A_j$, for every $i$, $d_{A_j}(i) \leq am/p_j$ for $a > 1/\ln(1/\delta)$ and $0 < \delta \leq 1/p_j$. Then*

$$Pr[L^j > 3\ln(1/\delta)am/p_j] \leq p_j\delta.$$

*Proof.* First, notice that we have the following equivalent expression for $L^j(b_j)$, $b_j \in [p_j]$:

$$L^j(b_j) = \sum_{i \in [n]} d_{A_j}(i) \cdot X_{h_j(i)=b_j}$$

where $X_{h_j(i)=b_j}$ denotes the indicator variable for the event that the hash function $h_j$ sends node $i$ to row a bin with $j$-th coordinate $b_j$. Since our hash function is perfectly random, we have a weighted balls-into-bins problem, where each ball is a node $i$ of weight $d_{A_j}(i)$. The total weight of the balls is $m$ and $d_{A_j}(i) \leq am/p_j$ for every $i \in [n]$; since we have $p_j$ bins, we can apply Lemma C.1 to obtain the probability bound. □

**Lemma B.5.** *Let $m \geq p^2$. If for every set of attributes $S \subseteq [r]$ and vector of values $I = (i_j)_{j \in S} \in [n]^S$, $d_S(I) \leq am/\prod_{j \in S} p_j$ where $a \geq 1/\ln(1/\delta))$ and $0 < \delta < 1$, then*

$$Pr\left[L > [3\ln(1/\delta)a]^r \frac{m}{p}\right] < 2^{r+1}p^2\delta/(3\ln(1/\delta)a^2).$$

*Proof.* The proof uses a recursive argument. For each $b_1 \in [p_1]$, we look at the tuples that have been hashed to all the bins $b'$ such that $b'_1 = b_1$. These tuples form a sub-relation of $R$, which we denote by $R^{b_1}$; for $S \subseteq [r] - \{1\}$, and $I = (i_j)_{j \in S}$, let

$$d_S^{b_1}(I) = \sum_{i_1,\ h_1(i_1)=b_1} d_{S \cup \{1\}}(i_1, I)$$

denote the induced degrees of partial vectors of values in $R^{b_1}$.

We will show next that, with some high probability, our problem reduces to hashing the relations $R^{b_1}$ over $p_2 \times \cdots \times p_r$ bins such that for each such relation:

- $|R^{b_1}| \leq 3\ln(1/\delta)am/p_1$ which we denote by $m_1$, and

- for each subset of attributes $S \subseteq [r] - \{1\}$ and vector of values $I = (i_j)_{j \in S} \in [n]^S$, we have that $d_S^{b_1}(I) \leq am_1 / \prod_{j \in S} p_j$.

The first task it to get a bound on the size of $R^{b_1}$. Since $|R^{b_1}| = L^1(b_1)$, we can apply Lemma B.4 to obtain that $Pr[\max_{b_1}\{|R^{b_1}|\} > 3\ln(1/\delta)am/p_1] \leq p_1\delta$.

The next task is to bound the degrees of vectors of vertices in each subrelation $R^{b_1}$ so that we can apply a recursive argument. In particular, for each $S \subseteq [r] - \{1\}$ and $I = (i_j)_{j \in S}$ we will show that with high probability, $d_S^{b_1}(I) \leq am_1 / \prod_{j \in S} p_j$.

Fix some $S \subseteq [r] - \{1\}$. For $I = (i_j)_{j \in S} \in [n]^S$, if $d_S(I) \leq am_1 / \prod_{j \in S} p_j = 3\ln(1/\delta)a^2 m/(p_1 \prod_{j \in S} p_j)$ then we automatically have $d_S^{b_1}(I) \leq am_1 / \prod_{j \in S} p_j$ for every $b_1 \in [p_1]$. Since $|R| \leq m$, there are fewer than $p_1 \prod_{j \in S} p_j / (3a^2 \ln(1/\delta))$ choices $I \in [n]^S$ such that

$$d_S(I) > am_1 / \prod_{j \in S} p_j = 3\ln(1/\delta)a^2 m/(p_1 \prod_{j \in S} p_j).$$

We are interested in bounding the load from each of these choices on a given $R_{b_1}$. By our hypothesis on $R$, for each such $I$ there are $d_{S \cup \{1\}}(i_1, I) \leq am/(p_1 \prod_{j \in S} p_j)$ elements of $R$ that have any fixed set of values $(i_1, I) = (i_j)_{j \in S \cup \{1\}}$. For each of the choices of $I$ the load is given by the load of a balls-in-bins problem with $p_1$ bins, total weight $d_S(I) \leq am/\prod_{j \in S} p_j$, and maximum weight per ball of $\max_{i_1} d_{S \cup \{1\}}(i_1, I) \leq am/(p_1 \prod_{j \in S} p_j)$. By Lemma C.1 the maximum value of $d_S^{b_1}(I)$ for each $b_1$ is at most $3\ln(1/\delta)am/(p_1 \prod_{j \in S} p_j) = am_1/\prod_{j \in S} p_j$ except with probability at most $p_1\delta$. For this value of $S$, there are at most $p_1 \prod_{j \in S} p_j / (3a^2 \ln(1/\delta))$ choices of $i_j$ to which we need to apply this argument. Running over all these choices of $S$, we get a total failure probability for the reduction step of at most

$$[p_1 + p_1^2 \sum_{S \subseteq [r] - \{1\}} \prod_{j \in S} p_j / (3a^2 \ln(1/\delta))]\delta,$$

which we denote by $P(m, p_1, \ldots, p_r)$.

We now have the precisely analogous conditions for each of the $p_1$ choices of $R^{b_1}$ as we had initially for $R$, except with $m_1 = 3\ln(1/\delta)am/p_1$ replacing $m$ and the first coordinate removed. Let $K = 3\ln(1/\delta)a$. Repeating this $j \leq r$ times, for each $j$ if all the reduction steps succeed we have total load at most $m_j = K^j m / \prod_{i \leq j} p_i$ in the bins consistent with $(b_1, b_2, \ldots, b_j)$ and the total failure probability for all $r$ levels is at most

$$\begin{aligned}
&\delta \cdot \sum_{j=1}^{r} P((K^j m / \prod_{i<j} p_i), p_j, \ldots, p_r) \prod_{i<j} p_i \\
&\leq \delta \cdot \left[ \sum_{j=1}^{r} (\prod_{i=1}^{j} p_i)[1 + \sum_{S \subseteq [r],\, \min(S)=j} (\prod_{i \in S} p_i)/(aK)] \right] \\
&\leq 2\delta \cdot \sum_{j=1}^{r} [\sum_{S \subseteq [r],\, \min(S)=j} p_j p/(aK)] \\
&< 2^{r+1}\delta p^2/(aK).
\end{aligned}$$

which is what we wanted to prove. □

As a corollary, we have:

**Corollary B.6.** *If for every subset of attributes $S \subseteq [r]$ and vector of values $I = (i_j)_{j \in S}$ we have that $d_S(I) \leq am/\prod_{j \in S} p_j$, then the probability that the load $L > [9a \ln p]^r m/p$ is polynomially small in $p$.*

*Proof.* Let $\delta = p^{-3}$ so $\ln(1/\delta) = 3 \ln p$ and apply Lemma B.5. □

**Proposition B.7.** *For any relation $R$, we have:*

$$Pr\left[L > \frac{(3r+1)m}{\min_j(p_j)}\right] \leq 2\sum_j p_j e^{-\left(m/\min_j(p_j)\right)^{1/r}/3}.$$

*Proof.* For attribute $A_j$, let us fix a threshold $\tau_j = m/(p_j \ln(1/\delta))$, where $\delta$ is a parameter we will choose later. Let $R^j$ denote the subset of $R$ that contains tuples $t$ such that $d_j(t[A_j]) \leq \tau_j$.

If $L_0^j$ denotes the load $L^j$ for the instance $R^j$, assuming the worst-case scenario where $|R_j| = m$, we can apply directly Lemma B.4 to obtain that $Pr[L_0^j > 3m/p_j] \leq p_j\delta$. Notice also that the maximum load of any bin for $R^j$ is upper bounded by $L_0^j$.

Next consider the tuples in $R' = R \setminus \bigcup_j R^j$, which are the tuples $t$ such that for every attribute $A_j$, we have $d_j(t[A_j]) > \tau_j$. The main observation here is that the tuples in $R'$ are very few. Indeed for each $A_j$ there are at most $m/\tau_j = p_j \ln(1/\delta)$ values of the $A_j$ attribute that appear in tuples of $R'$ and hence

$$|R'| \leq \prod_{i=1}^{r}(p_j \ln(1/\delta)) = p \ln^r(1/\delta).$$

Since only $p_j \ln(1/\delta)$ values of the $A_j$ attribute appear in $R'$ and there are $p_j$ values of the bin coordinate for $A_j$, Corollary C.2 implies that the probability that some bin will receive $> 3\ln(1/\delta)$ values from the attribute $A_j$ is at most $p_j\delta$. Hence, by applying a union bound, the probability that there exists a bin that receives $> 3\ln(1/\delta)$ values on any coordinate is at most $\delta \sum_j p_j$. Notice that if this bad event does not happen, each bin will hold at most $[3\ln(1/\delta)]^r$ tuples from $R'$.

Since $R = R' \cup (\bigcup_j R_j)$, we can add up the probabilities of the bad events to conclude that

$$\Pr\left[L > \frac{3rm}{\min_j(p_j)} + [3\ln(1/\delta)]^r\right] \leq 2\delta\sum_j p_j$$

The claim is obtained by choosing the value of $\delta$ such that $[3\ln(1/\delta)]^r = m/\min_j(p_j)$; that is $\delta = \exp(-[m/\min_j(p_j)]^{1/r}/3)$. □

We can easily now obtain the following corollary.

**Corollary B.8.** *If $m \geq p \ln^c(p)$ for some large constant $c$, then $L > \frac{(3r+1)m}{\min_j p_j}$ with probability polynomially small in $p$.*

# C Balls Into Bins

Here, we study the standard setting where balls (possibly weighted) are thrown independently and uniformly at random into bins, and we are interested in the maximum weight that each bin will receive.

**Lemma C.1.** *Let $0 < \delta < 1$. Suppose that we throw weighted balls independently and uniformly at random into $p$ bins, such that:*

1. *The total weight of the balls is at most $m$.*
2. *The weight of every ball at most $B = am/p$, for $a \geq 1/\ln(1/\delta)$.*

*Then, the probability that the maximum weight per bin is $> 3\ln(1/\delta)am/p$ is at most $p\delta$.*

*Proof.* By convexity, the maximum load is maximized for the weight distribution that has $n = m/B$ balls of weight $B$; hence, we will upper bound the probability of this event. Consider some bin $j = 1, \ldots, p$. Each ball falls into this bin independently with probability $1/p$. If $X_{ij}$ denotes the event that the ball $i = 1, \ldots, n$ falls into the bin $j$, the random variable $W_j = B\sum_i X_{ij}$ denotes the weight of the bin $j$. We now apply a form of Chernoff bound to the random variable $X_j = \sum_i X_{ij}$ which has $\mu = E[X_j] = m/(pB) = 1/a$ which says that

$$Pr[X_j > (1+\epsilon)\mu] \leq \begin{cases} \exp(-\mu\epsilon^2/4) & \epsilon < 2e-1 \\ 2^{-\epsilon\mu} & \epsilon \geq 2e-1. \end{cases}$$

We set $\epsilon = 2a\ln(1/\delta)$ and observe that $\epsilon \leq 3a\ln(1/\delta) - 1 \geq 2$ since $a \geq 1/\ln(1/\delta)$. We have two cases. For $2 \leq \epsilon < 2e-1$ we use the first bound to obtain $\Pr[X_j > 3\ln(1/\delta)] \leq \exp(-2\ln(1/\delta)\epsilon^2/4) \leq \exp(-2\ln(1/\delta)) = \delta^2 < \delta$; for $2a\ln(1/\delta) \geq 2e-1$ we use the second bound to obtain $\Pr[X_j > 3\ln(1/\delta)] \leq 2^{-\epsilon\mu} \leq 2^{-2\ln(1/\delta)} < \delta$. Applying a union bound we have $\Pr[\max_j W_j > 3\ln(1/\delta)am/p] \leq p\delta$ as required. □

**Corollary C.2.** *Suppose that we throw $m$ balls independently and uniformly at random into $p$ bins. Then, the probability that the maximum number of balls per bin is $> 3m/p$ is at most $pe^{-m/p}$.*

*Proof.* We apply Lemma C.1 with $\delta = e^{-m/p}$, $a = 1/\ln(1/delta) = p/m$, and hence $B = 1$. □

Notice that, in order to get any meaningful probability from Corollary C.2, it must be that $m \geq p\ln(p)$.

# D Proofs

For convenience of notation, for each relation put all tuples in a bin associated with the empty set of variables and give it bin exponent 0. Define an ordering on bin combinations $\mathcal{B}' = (\mathbf{x}', (\beta'_j)_j)$ and $\mathcal{B} = (\mathbf{x}, (\beta_j)_j)$ by $\mathcal{B}' < \mathcal{B}$ iff $\mathbf{x}' \subset \mathbf{x}$ and for all $j$, $\beta'_j \leq \beta_j$. Observe that for any tuple $\mathbf{h} \in C(\mathcal{B})$ and every $\mathbf{x}' \subset \mathbf{x}$ there is a unique bin combination $\mathcal{B}'$ on $\mathbf{x}'$ with $\mathcal{B}' < \mathcal{B}$ such that the projection of $\mathbf{h}$ on $\mathbf{x}'$, $\mathbf{h}'$ is in $C(\mathcal{B}')$.

**Lemma D.1.** *Let $\mathcal{B}' = (\mathbf{x}', (\beta'_j)_j) < \mathcal{B} = (\mathbf{x}, (\beta_j)_j)$ be bin combinations and assume that $\log_p C'(\mathcal{B}') = \alpha \leq 1$. If $\mathbf{x} - \mathbf{x}' = \mathbf{x}_j - \mathbf{x}'_j$ for some $j \in [\ell]$ (that is $\mathbf{x}$ only adds new variables from relation $S_j$ to $\mathbf{x}'$) then for every $\mathbf{h}' \in C'(\mathcal{B}')$ there are at most $p^{1-\alpha}/N_{\text{bc}}$ assignments $\mathbf{h}''$ such that $\mathbf{h}_j = \mathbf{h}'_j\mathbf{h}'' \in C'(\mathcal{B})$ is an overweight heavy hitter of $S_j$ for $\mathcal{B}'$.*

*Proof.* By definition $\mathbf{h}'_j$ is contained in at most $m_j/p^{\beta'_j}$ tuples of $S_j$. However, each of the assignments $\mathbf{h}_j = \mathbf{h}'_h\mathbf{h}''$ that is overweight for $\mathcal{B}'$ is contained in at least $N_{\text{bc}} \cdot m_j/(p^{\beta'_j + \sum_{i \in \mathbf{x}_j - \mathbf{x}'_j} e_i^{(\mathcal{B}')}})$ tuples of $S_j$ because it is overweight. Therefore the number of different $\mathbf{h}''$ is at most $p^{\sum_{i \in \mathbf{x}_j - \mathbf{x}'_j} e_i^{(\mathcal{B}')}}/N_{\text{bc}}$ which is $\leq p^{1-\alpha}/N$ by the properties of the LP for $\mathcal{B}'$. □

We can now prove Lemma 4.2

*Proof of Lemma 4.2.* The proof is by induction on the partial order of bin combinations. By definition $|C'(\mathcal{B}_\emptyset)| = |C(\mathcal{B}_\emptyset)| = 1$. Now consider some arbitrary bin combination $\mathcal{B} = (\mathbf{x}, (\beta_j)_j)$ and assume by induction that $C'(\mathcal{B}') \leq p$ for all bin combinations $\mathcal{B}' = (\mathbf{x}', (\beta'_j)_j) < \mathcal{B}$. By definition, for every $\mathbf{h} \in C(\mathcal{B})$ that is in $C'(\mathcal{B})$ there must be some bin combination $\mathcal{B}' = (\mathbf{x}', (\beta'_j)_j) < \mathcal{B}$, such that the $\mathbf{h}'$, the restriction of $\mathbf{h}$ to $\mathbf{x}'$, is in $C'(\mathcal{B}')$ and there is some $j$ such that $\mathbf{x} - \mathbf{x}' = \mathbf{x}_j - \mathbf{x}'_j \subset vars(S_j)$ and $\mathbf{h}_j$ is an overweight heavy hitter of $S_j$ for bin combination $\mathcal{B}'$. Let $\alpha = \log_p |C'(\mathcal{B}')|$. Applying Lemma D.1 we see that each of the $p^\alpha$ assignments $\mathbf{h}' \in C'(\mathcal{B}')$ has at most $p^{1-\alpha}/N_{\text{bc}}$ extensions $\mathbf{h}$ to $\mathbf{x}$ such that $\mathbf{h}$ is overweight. Therefore each $\mathcal{B}' < \mathcal{B}$ can contribute at most $p/N_{\text{bc}}$ elements of $C'(\mathcal{B})$. Since there are at most $N_{\text{bc}}$ choices of $\mathcal{B}'$, the lemma follows. □